MA4233 Dynamical Systems

Thang Pang Ern

These notes are based off **Prof. Wei Daren's** MA4233 Dynamical Systems materials.

This set of notes was last updated on May 30, 2026. If you would like to contribute a nice discussion to the notes or point out a typo, please send me an email at `thangpangern@u.nus.edu`.

Contents

Chapter 1

Introduction

1.1 Dynamics

From the point of view of Galilei and Newton, nature obeys unchanging laws that mathematics can describe. Such elegant observations are summarized by the laws of mechanics, where things behave and evolve in a way determined by fixed rules.

The key idea behind the Newtonian revolution is that the principles of nature can be expressed in terms of mathematics, and physical events can be predicted precisely. Similar ideas were held by Laplace, namely that determinism is at the heart of dynamical systems. However, this is not generally true. It took almost 100 years for people to understand the impossibility of prediction in dynamical systems. It was Poincaré who proposed to study approximate behaviour instead of exact solutions for dynamical systems[1].

Poincaré's approach concentrates on the study of long-term asymptotic behaviour, which requires direct methods that do not rely on prior explicit calculation of solutions. Since then, geometric methods and probabilistic phenomena have played a central role in modern dynamics.

Example 1.1 (antipodal rabbits). The first example we wish to study is the number of antipodal rabbits, whose number is very huge in Australia. Suppose the number of rabbits is x, and the time variable is t. We try to understand the way of describing the number of rabbits as a function $x(t)$ of time.

There are two important features related to the number of rabbits: they eat and reproduce. To simplify the model, we suppose that during any given period Δt, a fixed percentage of the number of rabbits will be born and a fixed percentage will die. Hence the increment $x(t+\Delta t)-x(t)$ is proportional to $x(t)\Delta t$. By taking $\Delta t \to 0$, we obtain

$$\frac{dx}{dt} = kx,$$

[1] The famous butterfly effect statement of Edward Lorenz is another example of such phenomena: a butterfly may flutter by in Rio and thereby cause a typhoon in Tokyo a week later.

where k is the relative growth rate of the number of rabbits, independent of time t. The differential equation can also be written as

$$\frac{dx}{x} = kdt.$$

By integrating both sides and performing some algebraic manipulation, we see that $x(t) = x(0)e^{kt}$.

Example 1.2 (the leaning rabbits of Pisa). In the year 1202, Leonardo of Pisa considered the following more real question of the number of rabbits compared with Example 1.1. The main difference is that the number of rabbits is relatively small. The question is as follows:

> A man has one pair of rabbits at a certain place entirely surrounded by a wall. We wish to know how many pairs can be bred from it in one year, if the nature of these rabbits is such that they breed every month one other pair and begin to breed in the second month after their birth. Let the first pair breed a pair in the first month, then duplicate it and there will be 2 pairs in a month. How many pairs of rabbits can be bred from one pair in one year?

Let b_n be the number of rabbits at month n. Then it is clear that this sequence of numbers satisfies $b_{n+1} = b_n + b_{n-1}$ with the starting values $b_0 = b_1 = 1$. An interesting fact is that, as the son of Bonaccio, Leonardo of Pisa was known as *filius Bonacci*, or 'son of good nature'; Fibonacci for short. The above sequence is the famous Fibonacci sequence.

Now we would like to understand the connection between the Fibonacci model and the previous exponential growth model. Based on the exponential growth model, we should expect

$$b_{n+1} \approx ab_n$$

when b_n is sufficiently large, where a is independent of n. Suppose this is true. Then

$$a^2 b_n = ab_{n+1} = b_{n+2} = b_{n+1} + b_n = (a+1)b_n.$$

This gives $a^2 = a + 1$. The solutions to this quadratic equation give the corresponding eigenvalues of the Fibonacci sequence.

Example 1.3 (the butterfly model). Instead of a simple exponential growth model, we may think of some more complicated model. Suppose that our butterflies live with limited food supplies. This means that, by way of malnutrition or starvation, the number of butterflies may reduce when time passes.

Hence we use the following model:

$$f(x) = k(1 - \alpha x)x,$$

where x is the present number of butterflies.

The only modification to our model is adding a linear correction to the growth rate k. Here α represents the rate at which fertility is reduced through competition. In other words, we can say that $\frac{1}{\alpha}$ is the maximal possible number of butterflies. It is clear that if the number of butterflies is larger than $\frac{1}{\alpha}$, then the next year's population will decrease until the number is reduced to $\frac{1}{\alpha}$.

Another important feature is that if $\alpha x \ll 1$, then $1-\alpha x \approx 1$ and thus $f(x) \approx kx$. This is also intuitively true: the population is too small to suffer from competition for food as a large population would.

If we write our evolution law as $x_{i+1} = kx_i(1-\alpha x_i)$, then by changing variables $y_i = \alpha x_i$, we obtain

$$y_{i+1} = \alpha x_{i+1} = \alpha k x_i(1-\alpha x_i) = k y_i(1-y_i).$$

This gives the famous logistic equation

$$y_{i+1} = k y_i(1-y_i).$$

Example 1.4 (binary search). Suppose $f : [a,b] \to \mathbb{R}$ is a continuous function and $f(a) < 0 < f(b)$. We try to find a value $c \in (a,b)$ such that $f(c) = 0$. We will use binary search to find the root of the function.

The binary search procedure is as follows:

- **Case 1:** Let $z = \frac{a+b}{2}$. If $f(z) = 0$, then we have found the root.
- **Case 2:** If $f(z) > 0$, replace the interval $[a,b]$ by the interval $[a,z]$, and repeat the procedure on this interval.
- **Case 3:** If $f(z) < 0$, replace the interval $[a,b]$ by the interval $[z,b]$, and apply the procedure here.

This binary search produces a sequence of nested intervals, cutting the length in half at every step. This method is also reliable: it gives ever better approximations to the solution at a guaranteed rate, with an error term that can be calculated.

Example 1.5 (Newton-Raphson method). The Newton-Raphson method is a bit fancier than binary search and also much faster. However, this method requires that the function is differentiable.

We start with an educated guess x_0 of the solution. Then the equation of the tangent line to the graph of f gives us the x-intercept of the tangent line:

$$x_1 = F(x_0) = x_0 - \frac{f(x_0)}{f'(x_0)}.$$

This simple procedure can be applied repeatedly by iterating F, and gives a possible sequence of approximations of the solution.

However, some additional assumptions are needed to ensure that the method will succeed. The key observation of this method is that some initial choices provide situations where the asymptotic behaviour is complicated. An important development of this method to points in function spaces is the Kolmogorov-Arnold-Moser method, or KAM method, which is essential to study whether our solar system is stable.

Example 1.6 (closed geodesics). Imagine that a particle is moving freely on the surface of a sphere. Since there are no external forces and no friction, such a particle moves at a constant speed with no change of direction. It is quite clear that the particle always returns to the starting point periodically. In particular, there are infinitely many ways of travelling in a periodic fashion.

In the case where the sphere is not perfectly round, or is slightly dented, or even badly deformed, there is no obvious reason that a freely moving particle automatically returns home. To summarise the question:

> Are there still infinitely many ways, on a given deformed sphere, of moving freely and periodically?

One of the key features of free particle motion is that the path of motion is always the shortest connection between any two points on it that are not too far apart. This is not only true for the round sphere, but is also universally true. Such paths are called geodesics. Hence the above question is equivalent to:

> On any sphere, no matter how deformed, are there always infinitely many closed geodesics?

This is not an easy question and was solved only not so long ago.

Example 1.7 (first digits of the powers of 2). In this example, we study the pattern of first digits of the powers of 2.

The powers of 2 begin as follows:

2	2048	2097152	2147483648	2199023255552
4	4096	4194304	4294967296	4398046511104
8	8192	8388608	8589934592	8796093022208
16	16384	16777216	17179869184	17592186044416
32	32768	33554432	34359738368	35184372088832
64	65536	67108864	68719476736	70368744177664
128	131072	134217728	137438953472	140737488355328
256	262144	268435456	274877906944	281474976710656
512	524288	536870912	549755813888	562949953421312
1024	1048576	1073741824	1099511627776	1125899906842624

The first digits of the 50 entries are

2481361251
2481361251
2481361251
2481361251
2481371251

It seems that the pattern is close to being periodic, but a small change appears at the end. By doing more computations, similar phenomena will appear again and again.

By looking at the frequency of each digit, we may obtain some interesting estimates. More precisely, let $F_d(n)$ be the number of those powers 2^m, where $m = 1, \dots, n$, that begin with digit d. The frequency with which d appears as the first digit among the first n powers of 2 is $\frac{F_d(n)}{n}$. Later in this course, we will show that this limit exists and

$$\lim_{n\to\infty} \frac{F_d(n)}{n} = \log(d+1) - \log d.$$

Example 1.8 (last digits of polynomials). Let $x_n = n^2$, where $n \in \mathbb{N}$. It is not difficult to see that the last digits of x_n are periodic. If we define $x_n = n^2\sqrt{2}$, then the sequence of last digits of x_n is far from being periodic.

Let $P_n(d)$ be the number of times the last digit is d in the set $\left\{i^2\sqrt{2}\right\}_{i=0}^{n-1}$. We consider $\frac{P_n(d)}{n}$ for large n. We will show in this course that

$$\lim_{n\to\infty} \frac{P_n(d)}{n} = \frac{1}{10}.$$

Moreover, if $x_n = n^{2p/q}$, where $p, q \in \mathbb{N}$ and p, q are coprime, one can prove that the last digits of x_n are periodic.

Example 1.9 (cellular automata). In the 1980s, there was a very popular game called the Game of Life invented by Conway. It is a model of something that lives in a two-dimensional integer lattice.

Each cell is at a point of a fixed integer lattice and has two states: ‘present’ and ‘not there’, or equivalently 1 and 0. The rule of the game is that the number of cells changes in discrete time steps in a particular way. Each cell checks the states of some of its neighbours, up to some distance, and depending on all these states, changes its own state accordingly.

For example, it is possible that if all immediate neighbours are present, then the cell dies. The same may happen if there are no neighbours at all. The game was popular because from relatively simple rules, one can design intriguing patterns.

If the number of cells is finite, then there is not too much to say about the long-term behaviour of the system since it only has finitely many states. In particular, the corresponding system will eventually become periodic.

When there are infinitely many cells, there is no reason for this kind of cycling through the same patterns. Systems of this kind are called cellular automata.

To simplify the notation, we just look at the integer points on the line. Accordingly, a state of the system is a sequence, each entry of which has one of finitely many values.

Now we are able to define cellular automata mathematically. Let Ω_N be the space of sequences whose entries have values $0,\dots,N-1$. All cells have the same rule for their development. It is given by a function

$$f:\{0,\dots,N-1\}^{2n+1}\to\{0,\dots,N-1\}.$$

This maps $2n+1$-character-long strings of states $0,\dots,N-1$ to a state. It follows from the definition of f that the input consists of the states of all neighbours up to distance n in both directions. Thus, the evolution of the whole system is given by the map

$$\Phi:\Omega_N\to\Omega_N \quad \text{where} \quad (\Phi(w))_i=f(w_{i-n},\dots,w_{i+n}).$$

In the special case where $N=n=1$ and $f(x_{-1},x_0,x_1)=x_1$, this means that every individual just chooses to follow its right neighbour's lead, and can be understood as 'the wave'. The above is a general description of cellular automata, whose interest goes beyond the Game of Life. If we think of each of these sequences as a stream of data, then the map Φ transforms these data to a code, which is known as a sliding block code.

Such a class of codes is suitable for real-time streaming data encoding or decoding. The general class of dynamical systems whose states are given by sequences is called symbolic dynamics. The wave is actually our favourite transformation, which is also called the shift. Symbolic dynamics is one of the most important models in dynamical systems, and more details will be provided later in this course.

Chapter 2

Systems with Stable Asymptotic Behaviour

2.1 Systems with Stable Asymptotic Behaviour

In this chapter, we study dynamical systems whose long-term behaviour is stable. The main idea is that, although non-linear maps can be complicated, their behaviour near a fixed point can often be understood by looking at a suitable linear approximation.

We begin with the simplest case. Suppose $x_{i+1} = f(x_i) = kx_i$, where $k > 0$. Then $x_n = k^n x_0$. If $k > 1$, then $x_n \to +\infty$ as $n \to \infty$, provided $x_0 > 0$. If $0 < k < 1$, then $x_n \to 0$ as $n \to \infty$.

More generally, suppose $f(x) = kx + b$. If $k \neq 1$, then the fixed point satisfies $x = kx + b$ and hence, $x = \frac{b}{1-k}$. We change variables by setting $y = x - \frac{b}{1-k}$. Then,

$$y_{i+1} = x_{i+1} - \frac{b}{1-k} = kx_i + b - \frac{b}{1-k}.$$

Since

$$x_i = y_i + \frac{b}{1-k},$$

we obtain

$$y_{i+1} = k\left(y_i + \frac{b}{1-k}\right) + b - \frac{b}{1-k} = ky_i + \frac{kb}{1-k} + b - \frac{b}{1-k} = ky_i.$$

Thus the affine system can be reduced to the linear system $y_{i+1} = ky_i$.

Linear maps in higher dimensions may have more complicated behaviour of individual orbits. Although not all orbits have the same long-term behaviour, knowing a small number of orbits is often sufficient to understand the dynamics of all the others.

In fact, linear maps are quite rare in many applications. Non-linear maps are much more common. One of the key methods for studying non-linear maps is linearisation: we use a good linear approximation to study the non-linear dynamics near a fixed point.

Such a linear approximation is useful when the orbits of the non-linear map stay sufficiently close to the reference point.

Proposition 2.1. Suppose F is a differentiable map of the line and $F(b) = b$. If all orbits of the linearisation of F at b are asymptotic to b, then all orbits of F that start sufficiently near b are asymptotic to b as well.

Let $f : \mathbb{R}^n \to \mathbb{R}^m$ be a map with continuous partial derivatives. Write $f = (f_1, \ldots, f_m)$. At each point, we define the differential of f as the linear map represented by the matrix of partial derivatives

$$Df = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \cdots & \frac{\partial f_2}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \frac{\partial f_m}{\partial x_2} & \cdots & \frac{\partial f_m}{\partial x_n} \end{pmatrix}.$$

Let $\|\cdot\|$ be the norm of Euclidean space induced from the inner product. We define the operator norm of Df by

$$\|Df\| = \max_{\substack{v \neq 0 \\ v \in \mathbb{R}^n}} \frac{\|Df(v)\|}{\|v\|} = \max_{\|v\|=1} \|Df(v)\|.$$

This quantity measures the maximal stretching factor of the linear map Df.

2.2 Contractions and Perturbations

We now introduce the contraction principle (Theorem 2.1). This is one of the most important individual facts in analysis and dynamical systems. It appears in many places, such as the inverse function theorem, the implicit function theorem, the stability of periodic points, the existence and uniqueness of differential equations, and the stable manifold theorem.

A map f in this course usually means that its domain and range lie in the same space. In many cases, the range is also contained in the domain. The identity map is denoted by id, where $\text{id}(x) = x$.

Definition 2.1 (Lipschitz map and contraction). Let (X, d) be a metric space and let $S \subseteq X$. A map $f : S \to X$ is said to be Lipschitz continuous with Lipschitz constant λ, or λ-Lipschitz, if

$$d(f(x), f(y)) \leq \lambda d(x, y) \quad \text{for all } x, y \in S.$$

If $\lambda < 1$, then f is a contraction. If f is Lipschitz continuous, we define

$$\mathrm{Lip}(f) = \sup_{x \neq y} \frac{d(f(x), f(y))}{d(x,y)}.$$

Example 2.1. The function $f(x) = \sqrt{x}$ defines a contraction on $[1, \infty)$. To see this, notice that for $x \geq 1$ and $t \geq 0$, we have

$$\left(\sqrt{x} + \frac{t}{2}\right)^2 = x + \sqrt{x}t + \frac{t^2}{4} \geq x + t.$$

Hence,

$$\sqrt{x+t} \leq \sqrt{x} + \frac{t}{2} \quad \text{so} \quad \sqrt{x+t} - \sqrt{x} \leq \frac{t}{2}.$$

This shows that f is $\frac{1}{2}$-Lipschitz on $[1, \infty)$, and hence f is a contraction.

Definition 2.2 (orbit). Let $f : X \to X$ be a map and let $x \in X$. If f is not invertible, the orbit of x is

$$\mathrm{Orb}(f)(x) = \{x, f(x), f^2(x), \ldots, f^n(x), \ldots\}.$$

If f is invertible, the orbit of x is

$$\mathrm{Orb}(f)(x) = \{\ldots, f^{-n}(x), \ldots, f^{-1}(x), x, f(x), \ldots, f^n(x), \ldots\}.$$

Definition 2.3 (fixed point and periodic point). A fixed point is a point $x \in X$ such that $f(x) = x$. The set of fixed points is denoted by $\mathrm{Fix}(f)$.

A periodic point is a point $x \in X$ such that $f^n(x) = x$ for some $n \in \mathbb{N}$. Such an n is called a period of x. The smallest such n is called the prime period of x.

Example 2.2. Let $f(x) = -x^3$. Then 0 is the only fixed point. Indeed, $-x^3 = x$ implies $x(x^2+1) = 0$ and so $x = 0$. However, $f(1) = -1$ and $f(-1)$ so $\{-1, 1\}$ is a periodic orbit with prime period 2.

Definition 2.4 (exponential convergence). Let (X,d) be a metric space. We say that two sequences $\{x_n\}_{n \in \mathbb{N}}$ and $\{y_n\}_{n \in \mathbb{N}}$ in X converge exponentially to each other if there exist constants $c > 0$ and $\alpha \in (0,1)$ such that

$$d(x_n, y_n) < c\alpha^n.$$

In particular, if one of the sequences is constant, say $y_n = y$, then we say that x_n converges exponentially to y.

Theorem 2.1 (contraction principle). Let (X,d) be a complete metric space. Under the action of iterates of a contraction $f : X \to X$, all points converge with exponential speed to the unique fixed point of f.

Proof. Since f is a contraction, by Definition 2.1, there exists $\lambda \in (0,1)$ such that

$$d(f(x), f(y)) \leq \lambda d(x,y)$$

for all $x, y \in X$. By iteration, we obtain

$$d(f^n(x), f^n(y)) \leq \lambda^n d(x,y) \to 0$$

as $n \to \infty$. Hence the asymptotic behaviour of all points is the same.

Now fix $x \in X$. For $m \geq n$, we have

$$\begin{aligned} d(f^m(x), f^n(x)) &\leq \sum_{k=0}^{m-n-1} d(f^{n+k+1}(x), f^{n+k}(x)) \\ &\leq \sum_{k=0}^{m-n-1} \lambda^{n+k} d(f(x), x) \\ &\leq \frac{\lambda^n}{1-\lambda} d(f(x), x). \end{aligned}$$

Since $\lambda^n \to 0$, the sequence $\{f^n(x)\}_{n\in\mathbb{N}}$ is Cauchy. Since X is complete, the limit exists. Denote this limit by

$$x_0 = \lim_{n\to\infty} f^n(x).$$

Moreover, the first estimate shows that this limit is independent of the starting point x.

We now show that x_0 is a fixed point. For every $n \in \mathbb{N}$,

$$\begin{aligned} d(x_0, f(x_0)) &\leq d(x_0, f^n(x)) + d(f^n(x), f^{n+1}(x)) + d(f^{n+1}(x), f(x_0)) \\ &\leq d(x_0, f^n(x)) + \lambda^n d(x, f(x)) + \lambda d(f^n(x), x_0) \\ &= (1+\lambda) d(x_0, f^n(x)) + \lambda^n d(x, f(x)). \end{aligned}$$

Letting $n \to \infty$, the right hand side tends to 0. Hence $d(x_0, f(x_0)) = 0$ and therefore, $f(x_0,) = x_0$. Finally, if p and q are fixed points, then

$$d(p,q) = d(f(p), f(q)) \leq \lambda d(p,q).$$

Since $\lambda < 1$, this forces $d(p,q) = 0$, and hence $p = q$. Thus the fixed point is unique. □

The contraction principle is a global theorem. It assumes that the map is a contraction on the whole complete metric space. In many applications, however, a map is only contracting near a fixed point. This motivates the study of local contractions.

Definition 2.5 (convex set). A set $C \subseteq \mathbb{R}^n$ is convex if for all $a, b \in C$, the line segment with endpoints a and b is entirely contained in C. Equivalently, for all

$a, b \in C$ and $t \in [0,1]$,

$$(1-t)a + tb \in C.$$

Example 2.3. The disk $\{(x,y) \in \mathbb{R}^2 : x^2 + y^2 < 1\}$ is convex, whereas the annulus $\{(x,y) \in \mathbb{R}^2 : 1 < x^2 + y^2 < 2\}$ is not convex.

Theorem 2.2. If $C \subseteq \mathbb{R}^n$ is convex and open, and $f : C \to \mathbb{R}^m$ is differentiable with

$$\|Df(x)\| \leq M$$

for all $x \in C$, then for $x, y \in C$,

$$\|f(x) - f(y)\| \leq M\|x - y\|.$$

In particular, if $f : C \to C$ is a map with continuous partial derivatives and

$$\|Df(x)\| \leq \lambda < 1$$

at every point $x \in C$, then f is a λ-contraction.

Proof. Since C is convex, the line segment connecting x and y lies inside C. Let $c(t) = x + t(y-x)$, where $t \in [0,1]$. Define $g(t) = f(c(t))$. By the chain rule,

$$g'(t) = Df(c(t))(y-x). \tag{2.1}$$

Hence

$$\|f(y) - f(x)\| = \|g(1) - g(0)\| \leq \int_0^1 \|g'(t)\| \, dt$$

where we used the fundamental theorem of calculus. We then use (2.1) to show that this is

$$\int_0^1 \|Df(c(t))(y-x)\| \, dt \leq \int_0^1 M\|y-x\| \, dt = M\|y-x\|.$$

The final statement follows by taking $M = \lambda < 1$. □

Definition 2.6 (closed neighbourhood). A closed neighbourhood of a point x is the closure of an open set containing x.

Proposition 2.2. Let f be a continuously differentiable map with a fixed point x_0, where $\|Df_{x_0}\| < 1$. Then there is a closed neighbourhood U of x_0 such that $f(U) \subseteq U$ and f is a contraction on U.

Proof. Since Df is continuous, there is a small open ball $B(x_0, 2\eta)$ around x_0 on which $\|Df_x\| \leq \lambda < 1$. Let $U = \overline{B(x_0, \eta)}$. If $x, y \in U$, then the line segment joining x and y lies inside $B(x_0, 2\eta)$. Therefore, by Theorem 2.2,

$$d(f(x), f(y)) \leq \lambda d(x, y).$$

Hence f is a contraction on U.

Since $f(x_0) = x_0$, if $x \in U$, then

$$d(f(x), x_0) = d(f(x), f(x_0)) \le \lambda d(x, x_0) \le \lambda \eta < \eta.$$

Thus $f(x) \in U$. Hence $f(U) \subseteq U$. □

Proposition 2.3. Let f be a continuously differentiable map with a fixed point x_0, where $\|Df_{x_0}\| < 1$. Then there exists a closed neighbourhood U of x_0 satisfying $f(U) \subseteq U$ such that any map g sufficiently close to f is a contraction on U.

More precisely, if $\varepsilon > 0$, then there exists $\delta > 0$ such that any map g satisfying

$$\|g(x) - f(x)\| \le \delta \quad \text{and} \quad \|Dg_x - Df_x\| \le \delta$$

on U, maps U into itself and is a contraction on U, with its unique fixed point y_0 lying in $B(x_0, \varepsilon)$.

Proof. Since Df_x depends continuously on x, there is a small open ball $B(x_0, 2\eta)$ around x_0 on which $\|Df_x\| \le \lambda < 1$. Assume $\eta, \varepsilon < 1$, and take $\delta = \frac{\varepsilon\eta(1-\lambda)}{2}$. Then for every $x \in B(x_0, 2\eta)$,

$$\|Dg_x\| \le \|Dg_x - Df_x\| + \|Df_x\| \le \delta + \lambda \le \lambda + \frac{1-\lambda}{2} = \frac{1+\lambda}{2} < 1.$$

Hence g is a contraction on $B(x_0, 2\eta)$.

Let $U = \overline{B(x_0, \eta)}$. If $x \in U$, then

$$\begin{aligned} d(g(x), x_0) &\le d(g(x), g(x_0)) + d(g(x_0), f(x_0)) + d(f(x_0), x_0) \\ &\le \frac{1+\lambda}{2} d(x, x_0) + \delta + 0 \\ &\le \frac{1+\lambda}{2} \eta + \delta \end{aligned}$$

By taking δ sufficiently small, this is at most η. Hence $g(U) \subseteq U$. By the contraction principle (Theorem 2.1), g has a unique fixed point $y_0 \in U$. Since $g^n(x_0) \to y_0$, we have

$$d(x_0, y_0) \le \sum_{n=0}^{\infty} d(g^n(x_0), g^{n+1}(x_0)) \le d(g(x_0), x_0) \sum_{n=0}^{\infty} \left(\frac{1+\lambda}{2}\right)^n \le \frac{2\delta}{1-\lambda} = \varepsilon\eta.$$

Since $\eta < 1$, we obtain $d(x_0, y_0) < \varepsilon$. Thus, $y_0 \in B(x_0, \varepsilon)$. □

2.3 Attracting Fixed Points, Newton's Method, and Non-Decreasing Interval Maps

We now study fixed points more carefully. In particular, we distinguish between fixed points that merely keep nearby orbits nearby and fixed points that actually attract nearby orbits.

Example 2.4. Let the unit circle be represented by $\mathbb{R}/\mathbb{Z}$, and define

$$f(x) = x + \frac{\sin^2 x}{4}.$$

Then 0 is a semi-stable fixed point. This means that not every nearby orbit is asymptotic to 0.

Definition 2.7 (Poisson stable and asymptotically stable fixed points). A fixed point p is said to be Poisson stable if, for every $\varepsilon > 0$, there exists $\delta > 0$ such that if a point starts within δ of p, then its positive semi-orbit remains within ε of p (Figure 2.1a).

The point p is said to be asymptotically stable, or an attracting fixed point, if it is Poisson stable and there exists $a > 0$ such that every point within distance a of p is asymptotic to p (Figure 2.1b).

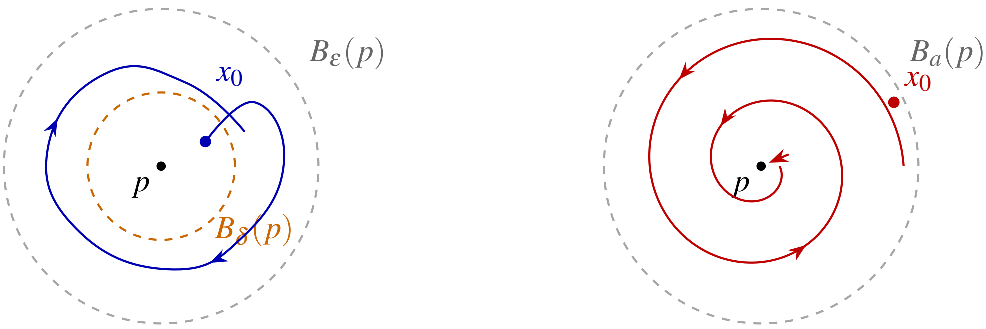

(a) Poisson stable fixed point   (b) Asymptotically stable fixed point

Figure 2.1: Comparison between Poisson stability and asymptotic stability. In the first case, the positive semi-orbit stays close to the fixed point p. In the second case, the positive semi-orbit stays close to p and converges to p as $t \to \infty$.

Note that Poisson stability does not imply asymptotic stability. For example, the system

$$\frac{dx}{dt} = y \quad \text{and} \quad \frac{dy}{dt} = -x$$

has many periodic orbits around the origin. Nearby orbits stay nearby, but they do not converge to the origin.

Lemma 2.1 (endpoint convergence for monotone interval maps). Let $I = [\alpha, \beta] \subseteq \mathbb{R}$ be a closed bounded interval and let $f : I \to I$ be a non-decreasing continuous map without fixed points in (α, β). Then one endpoint of I is fixed and all orbits converge to it, except possibly for the other endpoint if it is fixed as well.

If f is invertible and $f^{-1}(I) \subseteq I$, then both endpoints are fixed and all orbits of points in (α, β) are positively asymptotic to one endpoint and negatively asymptotic to the other.

Proof. Since $f(I) \subseteq I$, we have $f(\alpha) \geq \alpha$ and $f(\beta) \leq \beta$. Hence,

$$(f - \mathrm{id})(\alpha) \geq 0 \quad \text{and} \quad (f - \mathrm{id})(\beta) \leq 0.$$

By the intermediate value theorem, the continuous function $f - \mathrm{id}$ cannot change sign on I, because by assumption it has no zeros in (α, β).

Suppose first that $f(x) > x$ for every $x \in (\alpha, \beta)$. Then necessarily $f(\beta) = \beta$. For any $x \in (\alpha, \beta)$, define $x_n = f^n(x)$. Then $\{x_n\}$ is increasing and bounded above by β. By the monotone convergence theorem, it converges to some $x_0 \in (\alpha, \beta]$. By continuity of f,

$$f(x_0) = f\left(\lim_{n\to\infty} x_n\right) = \lim_{n\to\infty} f(x_n) = \lim_{n\to\infty} x_{n+1} = x_0.$$

Thus x_0 is a fixed point. Since there are no fixed points in (α, β), we must have $x_0 = \beta$. Therefore, $f^n(x) \to \beta$. The case $f(x) < x$ is similar and gives $f^n(x) \to \alpha$.

If f is invertible, then $z = f(y) > y$ implies $f^{-1}(z) = y < z$. Thus if $f(x) > x$ on (α, β), then $f^{-1}(x) < x$ on (α, β). Hence backward iterates converge to α, while forward iterates converge to β. The other case is similar. □

Definition 2.8 (heteroclinic and homoclinic points). Let $f : X \to X$ be an invertible map. If $x \in X$ satisfies

$$\lim_{n\to\infty} f^{-n}(x) = a \quad \text{and} \quad \lim_{n\to\infty} f^n(x) = b,$$

then x is said to be heteroclinic to a and b. If $a = b$, then x is said to be a homoclinic point of a.

Example 2.5 (heteroclinic points for an interval map). Let $X = [0,1]$ and let $f : X \to X$ where $f(x) = x^2$. Then, f is invertible with inverse $f^{-1}(x) = \sqrt{x}$. The fixed points of f are 0 and 1. For any $x \in (0,1)$, we have $f^n(x) = x^{2^n} \to 0$ as $n \to \infty$. On the other hand, $f^{-n}(x) = x^{1/2^n} \to 1$ as $n \to \infty$. Therefore every point $x \in (0,1)$ is heteroclinic to 1 and 0. In negative time, the orbit tends to 1, while in positive time, the orbit tends to 0.

Example 2.6 (a model homoclinic point). Let $p = (0,0)$ and define a two-sided sequence of points

$$x_n = \left(\frac{n}{n^2+1}, \frac{1}{n^2+1}\right) \quad \text{where } n \in \mathbb{Z}.$$

Let

$$X = \{p\} \cup \{x_n : n \in \mathbb{Z}\} \subseteq \mathbb{R}^2.$$

Define $f : X \to X$ by $f(p) = p$ where $f(x_n) = x_{n+1}$. Then f is invertible, with $f^{-1}(p) = p$ and $f^{-1}(x_n) = x_{n-1}$. Since $x_n \to p$ as $n \to \infty$ and also as $n \to -\infty$, we have, for any fixed $k \in \mathbb{Z}$,

$$f^n(x_k) = x_{k+n} \to p$$

as $n \to \infty$, and

$$f^{-n}(x_k) = x_{k-n} \to p$$

as $n \to \infty$. Hence each point x_k is a homoclinic point of p.

Proposition 2.4. Let $I \subseteq \mathbb{R}$ be a closed bounded interval and let $f : I \to I$ be a non-decreasing continuous map. Then every $x \in I$ is either fixed or asymptotic to a fixed point of f.

If f is increasing and $f^{-1}(I) \subseteq I$, then every $x \in I$ is either fixed or heteroclinic to adjacent fixed points.

Proof. The direction of motion is indicated by the sign of $f - \text{id}$. If $(f - \text{id})(x) < 0$, then $f(x) < x$ so x moves left. If $(f - \text{id})(x) > 0$, then $f(x) > x$ so x moves right.

We first show that fixed points exist. Write $I = [a,b]$. Since $f(I) \subseteq I$, we have $f(a) \geq a$ and $f(b) \leq b$. Hence,

$$(f - \text{id})(a) \geq 0 \quad \text{and} \quad (f - \text{id})(b) \leq 0.$$

By the intermediate value theorem, there exists $x \in I$ such that $(f - \text{id})(x) = 0$. Therefore $x \in \text{Fix}(f)$.

Since $\text{Fix}(f)$ is the set of zeros of the continuous function $f - \text{id}$, it is closed. If $\text{Fix}(f) = I$, then we are done.

Otherwise, $I \setminus \text{Fix}(f)$ is a nonempty open set and can be written as a disjoint union of open intervals. Let (α, β) be one of these intervals. Since f is non-decreasing, for every $y \in [\alpha, \beta]$,

$$\alpha = f(\alpha) \leq f(y) \leq f(\beta) = \beta.$$

Thus, $f([\alpha, \beta]) \subseteq [\alpha, \beta]$. By Lemma 2.1, every point in (α, β) converges to one of the endpoints, which is a fixed point.

If f is increasing and $f^{-1}(I) \subseteq I$, then the same argument applied to backward iterates shows that every non-fixed point is heteroclinic to adjacent fixed points. □

Definition 2.9 (repelling fixed point). A fixed point x is said to be a repelling fixed point, or a repeller, if there exists $\varepsilon_0 > 0$ such that for every $\varepsilon \in (0, \varepsilon_0)$ and every y within ε of x, there exists $n \in \mathbb{N}$ such that the positive semi-orbit of y eventually leaves the ε-neighbourhood of x (Figure 2.2).

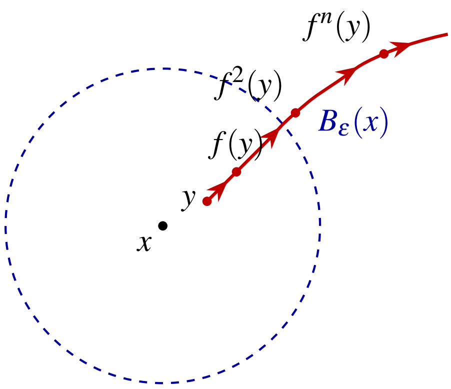

Figure 2.2: A repelling fixed point x. Any point y sufficiently close to x eventually has a positive iterate outside the ε-neighbourhood of x.

Definition 2.10 (topological transitivity and minimality). A homeomorphism $f : X \to X$ is said to be topologically transitive if there exists a point $x \in X$ such that its orbit

$$\mathrm{Orb}(f)(x) = \{f^n(x)\}_{n\in\mathbb{Z}}$$

is dense in X.

A homeomorphism $f : X \to X$ is said to be minimal if the orbit of every point $x \in X$ is dense in X.

Let $f : \mathbb{R} \to \mathbb{R}$ be a differentiable function. Recall that Newton's method helps us find a root of f quickly, provided we start with a reasonable initial guess. Suppose x_0 is our initial guess. The tangent line to the graph of f at $(x_0, f(x_0))$ has equation

$$y = f(x_0) + f'(x_0)(x - x_0).$$

The improved guess x_1 is the x-intercept of this tangent line. Setting $y = 0$, we obtain

$$0 = f(x_0) + f'(x_0)(x_1 - x_0).$$

Hence

$$x_1 = x_0 - \frac{f(x_0)}{f'(x_0)}.$$

Thus Newton's method is the iteration of the map

$$F(x) = x - \frac{f(x)}{f'(x)}.$$

Definition 2.11 (superattracting fixed point). A fixed point x of a differentiable map F is said to be superattracting if $F'(x) = 0$.

Proposition 2.5. If $|f'(x)| > \delta$ and $|f''(x)| < M$ on a neighbourhood of the root r, then r is a superattracting fixed point of the Newton map

$$F(x) = x - \frac{f(x)}{f'(x)}.$$

Proof. Since r is a root of f, we have $f(r) = 0$. Therefore,

$$F(r) = r - \frac{f(r)}{f'(r)} = r.$$

Thus r is a fixed point of F.

Differentiating

$$F(x) = x - \frac{f(x)}{f'(x)},$$

we get

$$F'(x) = 1 - \frac{f'(x)f'(x) - f(x)f''(x)}{(f'(x))^2} = \frac{f(x)f''(x)}{(f'(x))^2}.$$

Since $f(r) = 0$, it follows that $F'(r) = 0$. Hence r is a superattracting fixed point of F. □

Note that a small first derivative might cause the intersection of the tangent line with the x-axis to go very far from x_0. Hence Newton's method may fail if the initial choice is bad. The condition $|f''(x)| < M$ holds whenever f'' is continuous on a sufficiently small neighbourhood.

We now study the quadratic family

$$f_\lambda(x) = \lambda x(1-x).$$

Definition 2.12 (quadratic family). The family of maps

$$f_\lambda : [0,1] \to [0,1] \quad \text{where} \quad f_\lambda(x) = \lambda x(1-x),$$

where $\lambda \in [0,4]$, is called the quadratic family.

This is one of the most popular models in one-dimensional dynamics, both real and complex. For $\lambda < 3$, the corresponding dynamics are simple. For larger values of λ, the dynamics become much more complicated.

Proposition 2.6. For $0 \le \lambda \le 1$, all orbits of $f_\lambda(x) = \lambda x(1-x)$ on $[0,1]$ are asymptotic to 0.

Proof. For $x \neq 0$, we have

$$f_\lambda(x) = \lambda x(1-x) \le x(1-x) < x.$$

Thus $\{f_\lambda^n(x)\}_{n\in\mathbb{N}}$ is decreasing and bounded below by 0. Hence it converges.

The limit must be a fixed point of f_λ. Since the only fixed point in $[0,1]$ for $0 \le \lambda \le 1$ is 0, we get $f_\lambda^n(x) \to 0$. Moreover, if $0 \le \lambda < 1$, then

$$|f_\lambda'(x)| = \lambda|1-2x| \le \lambda < 1.$$

Hence f_λ is a contraction, and all orbits converge to 0 with exponential speed. □

Proposition 2.7. For $1 < \lambda \le 3$, all orbits of $f_\lambda(x) = \lambda x(1-x)$ on $[0,1]$, except for 0 and 1, are asymptotic to the fixed point $x_\lambda = 1-\lambda^{-1}$.

Proof. First, we identify the fixed points. Solving $f_\lambda(x) = x$, we get we get $\lambda x(1-x) = x$ so the fixed points are $x = 0$ and $x_\lambda = 1-\lambda^{-1}$. The second fixed point lies in $[0,1]$ precisely when $\lambda > 1$.

- **Case 1:** Suppose $1 < \lambda \le 2$. Then, $x_\lambda < \frac{1}{2}$. The map f_λ is increasing on $[0,x_\lambda]$, and $f_\lambda(x) > x$ for $x \in (0,x_\lambda)$. Hence every point of $[0,x_\lambda]$ is positively asymptotic to x_λ. Now notice that

  $$f_\lambda(1-x) = f_\lambda(x).$$

  Thus f_λ is symmetric about $x = \frac{1}{2}$, and hence $f_\lambda([1-x_\lambda,1]) \subseteq [0,x_\lambda]$. Therefore every point in $[1-x_\lambda,1)$ is also asymptotic to x_λ. It remains to consider the interval $[x_\lambda,1-x_\lambda]$. On this interval,

  $$|f_\lambda'(x)| = \lambda|1-2x| \le \lambda|1-2x_\lambda|.$$

  Since $x_\lambda = 1-\lambda^{-1}$, we have

  $$\lambda|1-2x_\lambda| = \lambda|1-2(1-\lambda^{-1})| = |2-\lambda| < 1.$$

  Thus f_λ is a contraction on $[x_\lambda,1-x_\lambda]$. Hence all points in this interval are asymptotic to the unique fixed point x_λ.

- **Case 2:** Suppose $2 < \lambda \le 3$. Then, $x_\lambda > \frac{1}{2}$ so f_λ is no longer increasing on $[0,x_\lambda]$. We use a more careful argument. Let

  $$I = \left[1-x_\lambda, f_\lambda\left(\frac{1}{2}\right)\right] = \left[\lambda^{-1}, \frac{\lambda}{4}\right].$$

  We first show that $f_\lambda(I) \subseteq I$. For every $x \in [0,1]$,

  $$f_\lambda(x) \le f_\lambda\left(\frac{1}{2}\right) = \frac{\lambda}{4}.$$

  For the lower bound, note that the minimum of f_λ on I occurs at $x = \frac{\lambda}{4}$. Thus

  $$f_\lambda\left(\frac{\lambda}{4}\right) = \frac{\lambda^2}{4} - \frac{\lambda^3}{16}.$$

  It is enough to show that

  $$\frac{\lambda^2}{4} - \frac{\lambda^3}{16} \ge \lambda^{-1}.$$

Equivalently,

$$\frac{\lambda^3}{4}\left(1-\frac{\lambda}{4}\right) \geq 1.$$

Let

$$q(\lambda) = \frac{\lambda^3}{4}\left(1-\frac{\lambda}{4}\right).$$

For $2 < \lambda < 3$, q is increasing, and $q(\lambda) > q(2) = 1$ so $f_\lambda(I) \subseteq I$.

Next, we show that every orbit except those of 0 and 1 eventually enters I. If $x \in I$, there is nothing to prove.

Suppose $x \in (0, \lambda^{-1})$. Let $x_n = f_\lambda^n(x)$. If the orbit never enters I, then $x_n \leq \lambda^{-1}$ for all n. But on $[0, \lambda^{-1}]$, we have $f_\lambda(x) > x$. Hence $\{x_n\}$ is increasing and bounded, so it converges to a fixed point in $(0, \lambda^{-1}]$. This is impossible, since there is no fixed point in this interval. Therefore the orbit must enter I.

Finally, for every $x \in \left[\frac{\lambda}{4}, 1\right]$, we have $f_\lambda(x) \in \left[0, \frac{\lambda}{4}\right)$. Thus these points also eventually enter I.

It remains to show that points in I converge to x_λ. Direct contraction may fail, so we study f_λ^2 instead of f_λ. Notice that

$$f_\lambda([\lambda^{-1}, x_\lambda]) \subseteq \left[x_\lambda, \frac{\lambda}{4}\right] \quad \text{and} \quad f_\lambda\left(\left[x_\lambda, \frac{\lambda}{4}\right]\right) \subseteq [\lambda^{-1}, x_\lambda].$$

Therefore,

$$f_\lambda^2([\lambda^{-1}, x_\lambda]) \subseteq [\lambda^{-1}, x_\lambda].$$

Now

$$f_\lambda^2\left(\frac{1}{2}\right) = \frac{\lambda^2(4-\lambda)}{16}.$$

By direct computation, this quantity remains in the appropriate range for $2 < \lambda \leq 3$. Let $J = \left[\frac{1}{2}, x_\lambda\right]$. Then, $f_\lambda^2(J) \subseteq \left(\frac{1}{2}, x_\lambda\right]$. Moreover, f_λ^2 is strictly increasing on J. Hence, by the result on non-decreasing interval maps, all points of J are asymptotic to a fixed point of f_λ^2. The only such fixed point is x_λ.

Since points in I eventually land in the relevant subintervals, it follows that every point in I converges to x_λ. Therefore every point in $[0,1]$, except 0 and 1, is asymptotic to $x_\lambda = 1 - \lambda^{-1}$.

This concludes the proof. The case distinction around $\lambda = 2$ appears because, at this value of the parameter, the fixed point x_λ has derivative zero. For larger values of λ, the derivative becomes negative, so nearby points approach the fixed point by alternating around it rather than approaching it monotonically.

Also, for $\lambda \geq 3$, the dynamics become much more complicated and will not be covered here. $\square$

Chapter 3

Circle Rotations

Recall that if our iteration f is a contraction, then every orbit is either fixed or attracted to a fixed point. However, this is not always the case. In general, non-trivial recurrence phenomena may occur.

We start with circle rotations to give some intuition about such phenomena.

Definition 3.1 (circle rotations). There are two convenient ways to represent the circle:

- **Multiplicative notation:** $S^1 = \{z \in \mathbb{C} : |z| = 1\} = \{e^{2\pi i\phi} : \phi \in \mathbb{R}\}$
- **Additive notation:** $S^1 = \mathbb{R}/\mathbb{Z}$

The map $e^{2\pi i\phi} \mapsto \phi$ identifies these two representations. We use R_α to denote rotation by angle $2\pi\alpha$. In multiplicative notation, $R_\alpha z = e^{2\pi i\alpha} z$; in additive notation, $R_\alpha x = x + \alpha \pmod 1$.

Under iteration, we have $R_\alpha^n z = e^{2\pi i n\alpha} z$ in multiplicative notation and $R_\alpha^n x = x + n\alpha \pmod 1$ in additive notation.

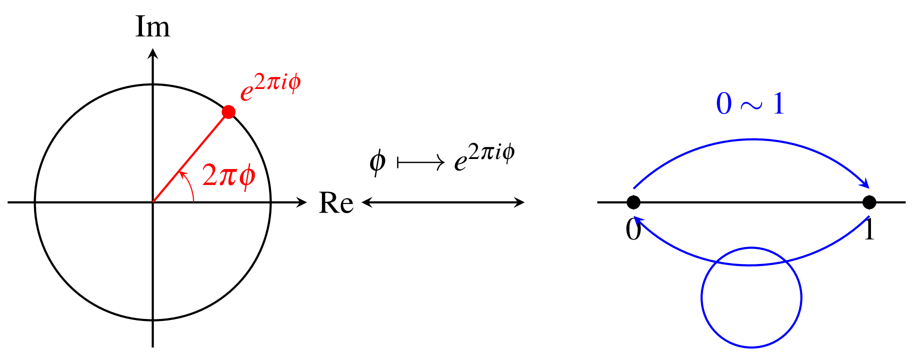

Multiplicative notation
$S^1 = \{z \in \mathbb{C} : |z| = 1\}$

Additive notation
$S^1 = \mathbb{R}/\mathbb{Z}$

If $\alpha = \frac{p}{q}$ for relatively prime integers p, q, then $R_\alpha^q = \mathrm{id}$. Hence the orbit of every point under a rational rotation is finite. Therefore, the interesting dynamics happen when $\alpha \notin \mathbb{Q}$.

3.1 Irrational Rotations and Density of Orbits

We now study the case where the rotation number α is irrational.

Proposition 3.1. If $\alpha \notin \mathbb{Q}$, then every positive semi-orbit of R_α is dense in S^1.

Proof. Let $x, z \in S^1$. We want to show that the positive orbit of x under R_α comes arbitrarily close to z.

Let $\varepsilon > 0$. Since $\alpha \notin \mathbb{Q}$, the positive semi-orbit of x is infinite. By the pigeonhole principle, among sufficiently many points in the orbit, two must be within distance ε of each other. More precisely, there exist $l, k \in \mathbb{N}$ with $l < k$ such that

$$d(R_\alpha^k(x), R_\alpha^l(x)) < \varepsilon.$$

Since R_α preserves distance, this gives

$$d(R_\alpha^{k-l}(x), x) < \varepsilon.$$

Let $\theta = (k-l)\alpha \pmod 1$, where $\theta \in [-1/2, 1/2]$. Then $R_\alpha^{k-l} = R_\theta$. Write $\rho = |\theta|$. Then $\rho < \varepsilon$.

Now the points $x, R_\theta(x), R_\theta^2(x), \ldots$ move around the circle by jumps of size ρ. Taking $N = \left\lfloor \frac{1}{\rho} \right\rfloor + 1$, the set $\{R_\theta^i(x) : i = 1, \ldots, N\}$ is ε-dense in S^1. Therefore, for every $z \in S^1$, there exists some $n \in \mathbb{N}$ such that

$$d(R_\alpha^n(x), z) < \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, the positive semi-orbit of x is dense in S^1. □

We can also prove the same result using invariant closed sets as follows.

Proof. Let $A \subseteq S^1$ be a non-empty closed invariant set. Suppose that $A \neq S^1$. Then $S^1 \setminus A$ is a non-empty open invariant set, so it is a union of disjoint open intervals.

Let I be one of the longest such intervals. Since rotations preserve length, all the intervals $R_\alpha^n(I)$ have the same length. Moreover, these intervals cannot overlap. If two of them overlap, then their union would contain an interval longer than I, contradicting the choice of I.

Since $\alpha \notin \mathbb{Q}$, no two iterates of I can coincide. Otherwise, some endpoint would return to itself, giving $k\alpha \in \mathbb{Z}$ for some non-zero k, which would imply $\alpha \in \mathbb{Q}$.

Thus we obtain infinitely many disjoint intervals of equal positive length inside S^1. This is impossible because the circle has finite length. Hence $A = S^1$. Therefore, every non-empty closed invariant set is the whole circle. It follows that every orbit is dense. □

Proposition 3.1 shows that irrational rotation is minimal. In particular, irrational rotation is also topologically transitive.

We now look at the frequencies with which iterates of a point visit various parts of the circle under an irrational rotation. Fix an arc $\Delta \subseteq S^1$. For $x \in S^1$ and $n \in \mathbb{N}$, let

$$F_\Delta(x,n) = \left|\{k \in \mathbb{Z} : 0 \le k < n,\ R_\alpha^k(x) \in \Delta\}\right|.$$

Thus $F_\Delta(x,n)$ counts the number of visits of the orbit segment $x, R_\alpha(x), \ldots, R_\alpha^{n-1}(x)$ to the arc Δ.

For fixed x and Δ, the function $F_\Delta(x,n)$ is non-decreasing in n. Since the positive orbit of every point is dense, we have $F_\Delta(x,n) \to \infty$ as $n \to \infty$. The important quantity is the relative frequency $\frac{F_\Delta(x,n)}{n}$.

Proposition 3.2. Suppose $\alpha \notin \mathbb{Q}$, and consider the rotation R_α. Let Δ and Δ' be arcs such that

$$l(\Delta) < l(\Delta').$$

Then there exists $N_0 \in \mathbb{N}$ such that if $x \in S^1$, $N \ge N_0$, and $n \in \mathbb{N}$, then

$$F_{\Delta'}(x, n+N) \ge F_\Delta(x,n).$$

Proof. Since the positive orbit of the left endpoint of Δ is dense, we can find $N_0 \in \mathbb{N}$ such that $R_\alpha^{N_0}(\Delta) \subseteq \Delta'$. Thus whenever $R_\alpha^n(x) \in \Delta$, we have $R_\alpha^{n+N_0}(x) \in \Delta'$. Therefore, for every $N \ge N_0$,

$$F_{\Delta'}(x, n+N) \ge F_{\Delta'}(x, n+N_0) \ge F_\Delta(x,n).$$

□

For arcs that are closed on the left and open on the right, we have an additivity property. If Δ_1 and Δ_2 are disjoint arcs that combine to form another arc, then

$$F_{\Delta_1}(x,n) + F_{\Delta_2}(x,n) = F_{\Delta_1 \cup \Delta_2}(x,n).$$

More generally, for any set A that is a finite union of disjoint arcs, we define

$$F_A(x,n) = \left|\{k \in \mathbb{Z} : 0 \le k < n,\ R_\alpha^k(x) \in A\}\right|.$$

Let

$$\overline{f}_x(A) = \limsup_{n\to\infty} \frac{F_A(x,n)}{n}.$$

Then $\overline{f}_x$ is subadditive, meaning that

$$\overline{f}_x(A_1 \cup A_2) \le \overline{f}_x(A_1) + \overline{f}_x(A_2).$$

Corollary 3.1. If $l(\Delta) < l(\Delta')$, then $\overline{f}_x(\Delta) \le \overline{f}_x(\Delta')$.

Similarly, define the lower asymptotic frequency by

$$\underline{f}_x(A) = \liminf_{n\to\infty} \frac{F_A(x,n)}{n}.$$

Since visits to A and visits to A^c are complementary, we have

$$\overline{f}_x(A) = 1 - \underline{f}_x(A^c).$$

Proposition 3.3. For any arc $\Delta \subseteq S^1$ and any $x \in S^1$,

$$\lim_{n\to\infty} \frac{F_\Delta(x,n)}{n} = l(\Delta).$$

Moreover, the convergence is uniform in x.

We first prove a lemma.

Lemma 3.1. If $l(\Delta) = \frac{1}{k}$, then

$$\overline{f}_x(\Delta) \le \frac{1}{k-1}.$$

Proof. Consider $k-1$ disjoint arcs $\Delta_1, \dots, \Delta_{k-1}$ of length $\frac{1}{k-1}$ each. Since $l(\Delta) < l(\Delta_i)$ for each i, the previous proposition gives a natural number N_i such that

$$F_{\Delta_i}(x, n+N_i) \ge F_\Delta(x,n).$$

Let $N = \max_i N_i$. Then,

$$F_{\Delta_i}(x, n+N) \ge F_\Delta(x,n)$$

for every i. Therefore,

$$(k-1)F_\Delta(x,n) \le \sum_{i=1}^{k-1} F_{\Delta_i}(x, n+N).$$

Since N is fixed, dividing by n and taking $n \to \infty$, we obtain

$$(k-1)\overline{f}_x(\Delta) \le 1.$$

Hence $\overline{f}_x(\Delta) \le \frac{1}{k-1}$. □

Corollary 3.2. If $l(\Delta) = \frac{p}{k}$ where $p, k \in \mathbb{N}$, then

$$\overline{f}_x(\Delta) \le \frac{p}{k-1}.$$

Proof. Split Δ into p arcs, each of length $\frac{1}{k}$. Applying the previous lemma to each arc and using subadditivity, we obtain $\overline{f}_x(\Delta) \le \frac{p}{k-1}$. □

We now discuss Birkhoff averages. Let A be a finite union of arcs. The characteristic function of A is

$$\chi_A(x) = \begin{cases} 1, & x \in A; \\ 0, & x \notin A. \end{cases}$$

Then

$$F_A(x,n) = \sum_{k=0}^{n-1} \chi_A(R_\alpha^k(x)).$$

In particular, for an arc Δ,

$$l(\Delta) = \int_{S^1} \chi_\Delta(x)\, dx.$$

Thus the previous proposition can be written as

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \chi_\Delta(R_\alpha^k(x)) = \int_{S^1} \chi_\Delta(x)\, dx.$$

Definition 3.2 (Birkhoff averaging operator). The Birkhoff averaging operator B_n associates to a function ϕ the function

$$B_n(\phi)(x) = \frac{1}{n} \sum_{k=0}^{n-1} \phi(R_\alpha^k(x)).$$

Proposition 3.4. The Birkhoff averaging operator has the following properties:

- B_n is linear: $B_n(a\phi + b\psi) = aB_n(\phi) + bB_n(\psi)$
- B_n is non-negative: if $\phi \geq 0$, then $B_n(\phi) \geq 0$
- B_n is non-expanding:

$$\sup_{x\in S^1} B_n(\phi)(x) \leq \sup_{x\in S^1} \phi(x)$$

- B_n preserves the average:

$$\int_{S^1} B_n(\phi)(x)\, dx = \int_{S^1} \phi(x)\, dx$$

Proposition 3.5. For any step function ϕ that is a linear combination of characteristic functions of arcs,

$$\lim_{n\to\infty} B_n(\phi) = \int_{S^1} \phi(x)\, dx.$$

Moreover, for any function ϕ that is a uniform limit of step functions, we also have

$$\lim_{n\to\infty} B_n(\phi) = \int_{S^1} \phi(x)\, dx.$$

Lemma 3.2. Every continuous function on S^1 is the uniform limit of step functions. The same is true for every function with finitely many discontinuities and one-sided limits at these discontinuities.

Proof. Every continuous function on S^1 is uniformly continuous. Thus, for every $\varepsilon > 0$, there exists $n \in \mathbb{N}$ such that on every arc of length $\frac{1}{n}$, the function varies by less than ε.

Divide S^1 into n arcs of equal length. Define a step function that is constant on each arc, taking for example the value of the original function at one chosen point of that arc. Then this step function differs from the original function by less than ε.

The same argument applies to functions with finitely many discontinuities and one-sided limits, after choosing the partition so that the discontinuities occur only at endpoints of the intervals. □

Proposition 3.6. If $\alpha \notin \mathbb{Q}$ and ϕ is continuous, then

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \phi(R_\alpha^k(x)) = \int_{S^1} \phi(y)\, dy$$

uniformly in x.

By approximation, we also obtain the following stronger version.

Proposition 3.7. If $\alpha \notin \mathbb{Q}$ and ϕ is Riemann integrable, then

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \phi(R_\alpha^k(x)) = \int_{S^1} \phi(y)\, dy$$

uniformly in x.

Proof. Pick a partition of S^1 into finitely many arcs I_i. The corresponding lower and upper Riemann sums can be interpreted as integrals of step functions.

Define $\phi_1 = \min \phi \mid_{I_i}$ on I_i and $\phi_2 = \max \phi \mid_{I_i}$ on I_i. Since ϕ is Riemann integrable, the partition can be chosen such that

$$\int_{S^1} \phi(x)\, dx - \varepsilon \leq \int_{S^1} \phi_1(x)\, dx \leq \int_{S^1} \phi_2(x)\, dx \leq \int_{S^1} \phi(x)\, dx + \varepsilon.$$

Since $\phi_1 \leq \phi \leq \phi_2$, we have

$$B_n(\phi_1) \leq B_n(\phi) \leq B_n(\phi_2).$$

Taking limits and using the result for step functions gives

$$\int_{S^1} \phi(x)\, dx - \varepsilon \leq \liminf_{n\to\infty} B_n(\phi) \leq \limsup_{n\to\infty} B_n(\phi) \leq \int_{S^1} \phi(x)\, dx + \varepsilon.$$

Letting $\varepsilon \to 0$, we obtain the desired result. □

Definition 3.3 (time average and space average). Given a function ϕ, we call

$$\lim_{n\to\infty}\frac{1}{n}\sum_{k=0}^{n-1}\phi(R_\alpha^k(x))$$

the time average of ϕ along the orbit of x.

The integral

$$\int_{S^1}\phi(y)\,dy$$

is called the space average of ϕ.

Definition 3.4 (uniquely ergodic). Let X be a compact metric space and let $f: X\to X$ be a continuous map. We say that f is uniquely ergodic if

$$\frac{1}{n}\sum_{k=0}^{n-1}\phi(f^k(x))$$

converges to a constant uniformly in x for every continuous function ϕ.

Thus irrational rotations are uniquely ergodic.

We now give another proof of the equality between time averages and space averages for irrational rotations.

Proof. Define the characters $c_m(x) = e^{2\pi imx}$. If $m\neq 0$, then

$$c_m(R_\alpha x) = e^{2\pi im(x+\alpha)} = e^{2\pi im\alpha}c_m(x).$$

Therefore,

$$\left|\frac{1}{n}\sum_{k=0}^{n-1}c_m(R_\alpha^k x)\right| = \left|\frac{1}{n}\sum_{k=0}^{n-1}e^{2\pi imk\alpha}\right| = \left|\frac{1-e^{2\pi imn\alpha}}{n(1-e^{2\pi im\alpha})}\right| \leq \frac{2}{n|1-e^{2\pi im\alpha}|}.$$

Since $\alpha\notin\mathbb{Q}$, we have $e^{2\pi im\alpha}\neq 1$ for $m\neq 0$. Hence,

$$\frac{1}{n}\sum_{k=0}^{n-1}c_m(R_\alpha^k x)\to 0.$$

By linearity, if

$$p(x) = \sum_{i=-l}^{l}a_i c_i(x)$$

is a trigonometric polynomial, then

$$\lim_{n\to\infty}B_n(p)(x)$$

exists and is constant. This constant must be

$$a_0 = \int_{S^1}p(x)\,dx,$$

because B_n preserves the integral.

Since continuous functions can be uniformly approximated by trigonometric polynomials, the same conclusion holds for all continuous functions. □

Note that if $f \in L^2(S^1)$, then its Fourier series converges to f in L^2. More precisely,

$$\lim_{n\to\infty} \int_{S^1} \left| f(x) - \sum_{k=-n}^{n} a_k c_k(x) \right|^2 dx = 0.$$

Next, we apply irrational rotations to the problem of first digits.

Proposition 3.8. Let $k \in \mathbb{N}$ be a natural number other than a power of 10, and let $p \in \mathbb{N}$. Then there exists $n \in \mathbb{N}$ such that p gives the initial digits of the decimal expansion of k^n.

Proof. The statement that p gives the initial digits of k^n means that for some $l \in \mathbb{N}$,

$$10^l p \le k^n < 10^l (p+1).$$

Taking logarithms base 10, this is equivalent to

$$l + \log p \le n \log k < l + \log(p+1).$$

Let $m = \lfloor \log p \rfloor + 1$. Then the condition can be written in terms of fractional parts as

$$\log\left(\frac{p}{10^{m-1}}\right) \le \{n \log k\} < \log\left(\frac{p+1}{10^{m-1}}\right).$$

Since k is not a power of 10, $\log k$ is irrational. Hence the sequence $\{n \log k\}_{n\in\mathbb{N}}$ is dense in $\mathbb{R}/\mathbb{Z}$. Therefore, for some n, the fractional part $\{n \log k\}$ lies in the required interval. Hence p appears as the initial digits of k^n. □

Proposition 3.9. Let $k \in \mathbb{N}$ not be a power of 10, and let $p \in \mathbb{N}$. Let $F_p^k(n)$ be the number of integers i between 0 and $n-1$ such that p gives the initial digits of the decimal expansion of k^i. Then

$$\lim_{n\to\infty} \frac{F_p^k(n)}{n} = \log(p+1) - \log p.$$

Moreover, this limit is independent of k.

Proof. The condition that p gives the initial digits of k^i is equivalent to the fractional part $\{i \log k\}$ lying in a fixed interval of length $\log(p+1) - \log p$. Since $\log k \notin \mathbb{Q}$, irrational rotation by $\log k$ is uniquely ergodic. Hence the frequency of visits to this interval is exactly its length. Therefore,

$$\lim_{n\to\infty} \frac{F_p^k(n)}{n} = \log(p+1) - \log(p).$$

□

3.2 Invertible Circle Maps

We now move from circle rotations to general invertible circle maps. The simple structure of the circle allows us to analyse the orbit structure of any invertible map of the circle. The two key ingredients are the ordering of points on the circle and the intermediate value theorem.

Although invertible circle maps are not as simple as circle rotations, we can define quantities that mimic circle rotations. The most important such quantity is the rotation number.

Recall that the circle can be written as $S^1 = \mathbb{R}/\mathbb{Z}$. Let $\pi : \mathbb{R} \to S^1$ be the projection map defined by $\pi(x) = \{x\}$, where $\{x\}$ is the equivalence class of x modulo $\mathbb{Z}$.

Proposition 3.10. If $f : S^1 \to S^1$ is continuous, then there exists a continuous map $F : \mathbb{R} \to \mathbb{R}$, called a lift of f, such that

$$f \circ \pi = \pi \circ F \quad \text{or equivalently} \quad f(\{x\}) = \{F(x)\}.$$

Such a lift is unique up to addition by an integer constant. Moreover,

$$\deg(f) = F(x+1) - F(x)$$

is an integer independent of x and independent of the choice of lift. This integer is called the degree of f.

If f is a homeomorphism, then $|\deg(f)| = 1$.

Proof. The idea is that the map f on the circle can be lifted continuously to a map F on the real line by choosing one starting value and then extending continuously.

Choose a point $p \in S^1$. Write $p = \{x_0\}$ for some $x_0 \in \mathbb{R}$, and write $f(p) = \{y_0\}$ for some $y_0 \in \mathbb{R}$. We define the lift F by requiring $F(x_0) = y_0$ and $f(\{x\}) = \{F(x)\}$. Continuity determines F locally, and repeating this construction extends F to all of $\mathbb{R}$.

Suppose $\widetilde{F}$ is another lift. Then for every $x \in \mathbb{R}$,

$$\{\widetilde{F}(x)\} = f(\{x\}) = \{F(x)\}.$$

Hence $\widetilde{F}(x) - F(x) \in \mathbb{Z}$. Since $\widetilde{F} - F$ is continuous and takes values in $\mathbb{Z}$, it must be constant. Therefore the lift is unique up to addition by an integer.

Also,

$$\{F(x+1)\} = f(\{x+1\}) = f(\{x\}) = \{F(x)\}.$$

Thus $F(x+1) - F(x) \in \mathbb{Z}$. Since this quantity depends continuously on x, it must be constant. This constant is called $\deg(f)$.

If f is a homeomorphism, then f must preserve or reverse the cyclic order of points. Therefore the degree must be either 1 or -1. □

Definition 3.5 (orientation preserving and orientation reversing). Suppose $f : S^1 \to S^1$ is invertible. If $\deg(f) = 1$, then f is called orientation-preserving. If $\deg(f) = -1$, then f is called orientation-reversing.

Notice that

$$F(x+1) - (x+1)\deg(f) = F(x) + \deg(f) - (x+1)\deg(f) = F(x) - x\deg(f).$$

Therefore $F(x) - x\deg(f)$ is periodic.

Lemma 3.3. If f is an orientation-preserving circle homeomorphism and F is a lift, then for all $x, y \in \mathbb{R}$,

$$F(y) - y \leq F(x) - x + 1.$$

Proof. Let $k = \lfloor y - x \rfloor$. Then, $0 \leq y - (x+k) < 1$. Now

$$\begin{aligned} F(y) - y &= F(y) - F(x+k) + F(x+k) - (x+k) + (x+k) - y \\ &= \big(F(x+k) - (x+k)\big) + \big(F(y) - F(x+k)\big) - \big(y - (x+k)\big) \end{aligned}$$

Since $F(x+1) = F(x) + 1$, we have $F(x+k) - (x+k) = F(x) - x$. Also, since F is increasing,

$$F(y) - F(x+k) \leq F(x+k+1) - F(x+k) = 1.$$

Therefore $F(y) - y \leq F(x) - x + 1$. □

We now introduce the rotation number. This quantity measures the average amount by which a lift moves points along the real line.

Proposition 3.11. Let $f : S^1 \to S^1$ be an orientation-preserving homeomorphism, and let $F : \mathbb{R} \to \mathbb{R}$ be a lift of f. Then

$$\rho(F) = \lim_{|n| \to \infty} \frac{1}{n}\left(F^n(x) - x\right)$$

exists for all $x \in \mathbb{R}$. The value $\rho(F)$ is independent of x and is well-defined up to an integer. More precisely, if $\widetilde{F}$ is another lift of f, then

$$\rho(F) - \rho(\widetilde{F}) = F - \widetilde{F} \in \mathbb{Z}.$$

Moreover, $\rho(F)$ is rational if and only if f has a periodic point.

Definition 3.6 (rotation number). Let $f : S^1 \to S^1$ be an orientation-preserving circle homeomorphism. The rotation number of f is

$$\rho(f) = \{\rho(F)\} \in \mathbb{R}/\mathbb{Z},$$

where F is any lift of f.

We first record a standard lemma from analysis.

Lemma 3.4. Suppose $\{a_n\}_{n\in\mathbb{N}}$ satisfies

$$a_{n+m} \leq a_n + a_m + L$$

for all $m, n \in \mathbb{N}$ and some constant L. Then

$$\lim_{n\to\infty} \frac{a_n}{n}$$

exists in $\mathbb{R} \cup \{-\infty\}$.

Proof. The condition says that the sequence is almost subadditive. Since $a_{n+1} \leq a_n + a_1 + L$, we get $a_n \leq na_1 + (n-1)L$. Thus $\liminf \frac{a_n}{n}$ is finite or equal to $-\infty$.

The usual subadditivity argument then shows that the lim sup and lim inf agree. Hence the limit exists. □

Definition 3.7 (factor, semiconjugacy, and conjugacy). Suppose $g : X \to X$ and $f : Y \to Y$ are maps of metric spaces. Suppose there exists a continuous surjective map $h : X \to Y$ such that $h \circ g = f \circ h$. Then f is said to be a factor of g, and h is called a semiconjugacy or factor map.

If h is a homeomorphism, then f and g are said to be conjugate, and h is called a conjugacy.

One important reason for introducing the rotation number is that it is a conjugacy invariant.

Proposition 3.12. If $f, h : S^1 \to S^1$ are orientation-preserving homeomorphisms, then

$$\rho(h^{-1} \circ f \circ h) = \rho(f).$$

Proof. Let F and H be lifts of f and h, respectively. Then

$$\pi \circ F = f \circ \pi \quad \text{and} \quad \pi \circ H = h \circ \pi.$$

Also, H^{-1} is a lift of h^{-1}. Therefore $H^{-1} \circ F \circ H$ is a lift of $h^{-1} \circ f \circ h$. Now

$$(H^{-1} \circ F \circ H)^n = H^{-1} \circ F^n \circ H.$$

Hence

$$\rho(H^{-1} \circ F \circ H) = \lim_{n\to\infty} \frac{H^{-1}(F^n(H(x))) - x}{n}.$$

Since H^{-1} differs from the identity by a bounded periodic function, this limit is the same as

$$\lim_{n\to\infty} \frac{F^n(H(x)) - H(x)}{n} = \rho(F).$$

Thus, $\rho(h^{-1} \circ f \circ h) = \rho(f)$. □

When the rotation number is rational, the circle homeomorphism has periodic points.

Proposition 3.13. Let $f : S^1 \to S^1$ be an orientation-preserving homeomorphism. Then all periodic orbits have the same period. More precisely, if

$$\rho(f) = \left\{ \frac{p}{q} \right\},$$

where $p, q \in \mathbb{Z}$ are relatively prime, and if F is a lift with $\rho(F) = \frac{p}{q}$, then $F^q(x) = x + p$ whenever $\{x\}$ is a periodic point.

Proof. Suppose $\{x\}$ is a periodic point. Then for some $r, s \in \mathbb{Z}$, $F^r(x) = x + s$. Therefore,

$$\frac{p}{q} = \rho(F) = \lim_{n\to\infty} \frac{F^{nr}(x) - x}{nr} = \frac{s}{r}.$$

Since p and q are relatively prime, there exists $m \neq 0$ such that $s = mp$ and $r = mq$. So, $F^{mq}(x) = x + mp$. We claim that $F^q(x) = x + p$. If $F^q(x) - p > x$, then by monotonicity, $F^{2q}(x) - 2p > x$ so by induction, $F^{mq}(x) - mp > x$ which contradicts $F^{mq}(x) = x + mp$.

Similarly, if $F^q(x) - p < x$, we obtain the opposite contradiction. Therefore, $F^q(x) = x + p$. □

This shows that, for rational rotation number, all periodic points are organised in a very rigid way.

Lemma 3.5. Let $I \subseteq \mathbb{R}$ be an interval whose endpoints are adjacent zeros of $F^q - \mathrm{id} - p$. Then, $F^q - \mathrm{id} - p$ has the same sign on the interiors of I and $F(I)$.

Proof. Suppose $F^q(x) - x - p > 0$ on I. Then, $F^q(x) > x + p$ for all $x \in I$. By monotonicity of F, for $x \in I$,

$$F^q(F(x)) = F(F^q(x)) > F(x + p) = F(x) + p.$$

Hence, $F^q - \mathrm{id} - p > 0$ on $F(I)$. The case where the sign is negative is similar. □

3.3 Circle Homeomorphisms Without Periodic Points

We finally consider orientation-preserving circle homeomorphisms with irrational rotation number. In this case, there are no periodic points.

Proposition 3.14. Let $F : \mathbb{R} \to \mathbb{R}$ be a lift of an orientation-preserving homeomorphism $f : S^1 \to S^1$ with $\rho(F) \notin \mathbb{Q}$. Then, for $n_1, n_2, m_1, m_2 \in \mathbb{Z}$ and $x \in \mathbb{R}$,

$$n_1\rho + m_1 < n_2\rho + m_2 \quad \text{if and only if} \quad F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2.$$

Proof. Since $\rho(F) \notin \mathbb{Q}$, the map f has no periodic points. Hence, $F^{n_1}(x) - F^{n_2}(x) \notin \mathbb{Z}$ for every x. Therefore the expression

$$F^{n_1}(x) + m_1 - F^{n_2}(x) - m_2$$

never vanishes. Since it is continuous in x, it cannot change sign. Thus the ordering is independent of x.

Suppose

$$F^{n_1}(x) + m_1 < F^{n_2}(x) + m_2.$$

Let $y = F^{n_2}(x)$. Then this is equivalent to $F^{n_1 - n_2}(y) - y < m_2 - m_1$ for all $y \in \mathbb{R}$.

Iterating this inequality gives

$$F^{r(n_1 - n_2)}(0) < r(m_2 - m_1)$$

for all $r \in \mathbb{N}$. Dividing by $r(n_1 - n_2)$ and letting $r \to \infty$, we obtain

$$\rho < \frac{m_2 - m_1}{n_1 - n_2}.$$

This is equivalent to $n_1\rho + m_1 < n_2\rho + m_2$. The converse follows by reversing the inequalities. $\square$

The previous proposition shows that, when the rotation number is irrational, the ordering of the orbit of a circle homeomorphism is the same as the ordering of the corresponding irrational rotation.

CHAPTER 4

Expanding Maps and Symbolic Dynamics

In this chapter, we study several important examples of chaotic dynamical systems. We begin with expanding maps on the circle, then study hyperbolic toral automorphisms, and finally introduce symbolic dynamical systems. The main theme is that complicated dynamics can often be understood by coding orbits using finitely many symbols.

This idea is extremely powerful. Instead of tracking an orbit exactly, we divide the phase space into finitely many pieces and record which piece the orbit visits at each time. The resulting sequence of symbols then gives a combinatorial model of the original dynamical system.

4.1 Expanding Maps

We now define expanding maps more generally.

Definition 4.1 (expanding map). A continuously differentiable map $f : S^1 \to S^1$ is said to be an expanding map if $|f'(x)| > 1$ for all $x \in S^1$.

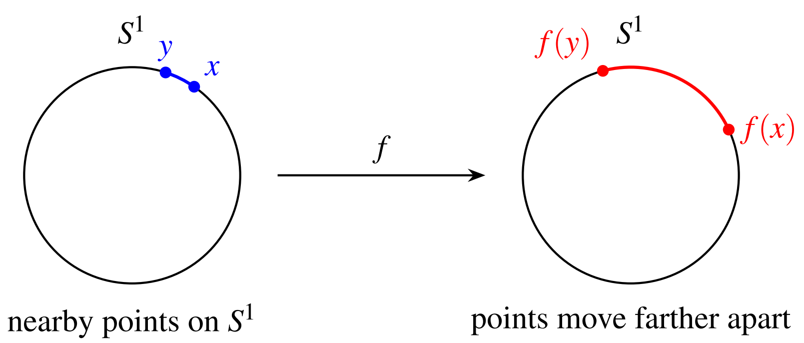

Since f is continuous, there exists a lift $F : \mathbb{R} \to \mathbb{R}$ such that $\{F(x)\} = f(\{x\})$ and $F(s+1) = F(s) + \deg(f)$, where $\deg(f)$ is the degree of f.

Lemma 4.1. If $f, g : S^1 \to S^1$ are continuous maps, then

$$\deg(g \circ f) = \deg(f)\deg(g).$$

In particular, $\deg(f^n) = \deg(f)^n$.

Proof. Let F and G be lifts of f and g, respectively. Since G is a lift of g, we have $G(s+1) = G(s) + \deg(g)$. More generally, for every integer k, we have

$$G(s+k) = G(s) + k\deg(g).$$

Now

$$F(s+1) = F(s) + \deg(f).$$

Therefore,

$$G(F(s+1)) = G(F(s) + \deg(f) = G(F(s)) + \deg(f)\deg(g).$$

Thus the lift $G \circ F$ of $g \circ f$ satisfies

$$(G \circ F)(s+1) = (G \circ F)(s) + \deg(f)\deg(g).$$

Hence, $\deg(g \circ f) = \deg(f)\deg(g)$. Taking $g = f^{n-1}$, we obtain $\deg(f^n) = \deg(f)^n$. □

Proposition 4.1. If $f : S^1 \to S^1$ is an expanding map, then $|\deg(f)| > 1$. Moreover, the number of fixed points of f^n is $P_n(f) = |\deg(f)^n - 1|$.

Proof. Since f is expanding on S^1, we have $|f'(x)| > 1$ for all $x \in S^1$. Passing to the lift $F : \mathbb{R} \to \mathbb{R}$, we also have $|F'(x)| > 1$. By the mean value theorem,

$$|\deg(f)| = |F(x+1) - F(x)| = |F'(y)|$$

for some $y \in [x, x+1]$. Hence, $|\deg(f)| > 1$.

By the previous lemma, $\deg(f^n) = \deg(f)^n$. Thus it is enough to count fixed points in the case $n = 1$, because fixed points of f^n are counted in the same way using the lift of f^n.

The fixed points of f are precisely the projections of points $x \in \mathbb{R}$ satisfying $F(x) - x \in \mathbb{Z}$. Define $g(x) = F(x) - x$. Then,

$$g(1) = F(1) - 1 = F(0) + \deg(f) - 1.$$

Therefore, $g(1) = g(0) + \deg(f) - 1$. By the intermediate value theorem, as x moves from 0 to 1, the function $g(x)$ crosses exactly $|\deg(f) - 1|$ integer values, except that if both endpoints are counted, the points 0 and 1 project to the same point of S^1.

Since $g'(x) = F'(x) - 1$ does not vanish in the relevant monotone case, g takes each integer value at most once. Hence there are exactly $|\deg(f) - 1|$ fixed points of f on S^1.

Applying this to f^n, we get

$$P_n(f) = |\deg(f^n) - 1| = |\deg(f)^n - 1|.$$

□

The proof also implies that for any continuous map, $P_n(f) \geq |\deg(f)^n - 1|$. For maps of degree 0, this merely guarantees a fixed point. For maps satisfying $|\deg(f)| > 1$ this gives exponential growth of the number of periodic points.

4.2 Hyperbolic Toral Automorphisms

Hyperbolic dynamics is one of the most important branches in modern dynamical systems. In this section, we study an elementary hyperbolic toral automorphism on $\mathbb{T}^2$, which provides intuition for the general theory.

Consider the linear map $L(x,y) = (2x+y, x+y)$. In matrix form,

$$L = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

Given two points (x,y) and (x',y') in $\mathbb{R}^2$ that represent the same point on $\mathbb{T}^2$, we have

$$(x - x', y - y') \in \mathbb{Z}^2.$$

Since L has integer entries, it follows that

$$L(x,y) - L(x',y') \in \mathbb{Z}^2.$$

Therefore $L(x,y)$ and $L(x',y')$ also represent the same element of $\mathbb{T}^2$. Hence L induces a well-defined map

$$F_L : \mathbb{T}^2 \to \mathbb{T}^2 \quad \text{where} \quad F_L(x,y) = (2x+y, x+y) \pmod 1.$$

This is an automorphism of the torus viewed as an additive group.

The matrix

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$

has determinant 1. Its inverse is

$$\begin{pmatrix} 1 & -1 \\ -1 & 2 \end{pmatrix},$$

which also has integer entries. Therefore F_L is invertible, and its inverse is induced by L^{-1}.

By direct computation, the eigenvalues of L are $\lambda_1 = \frac{3+\sqrt{5}}{2} > 1$ and $\lambda_2 = \frac{3-\sqrt{5}}{2} = \lambda_1^{-1} < 1$. Thus, one eigendirection is expanded and the other eigendirection is contracted. This is the source of hyperbolic behaviour.

Proposition 4.2. Periodic points of F_L are dense in $\mathbb{T}^2$, and

$$P_n(F_L) = \lambda_1^n + \lambda_1^{-n} - 2.$$

Proof. We first show that all points with rational coordinates are periodic points. Let $x = \frac{s}{q}$ and $y = \frac{t}{q}$, where $s,t,q \in \mathbb{Z}$. Then

$$F_L\left(\frac{s}{q},\frac{t}{q}\right) = \left(\frac{2s+t}{q},\frac{s+t}{q}\right) \pmod 1.$$

Thus the image again has rational coordinates with denominator q. There are only finitely many points on $\mathbb{T}^2$ whose coordinates can be represented with denominator q. Hence the forward orbit must eventually repeat:

$$F_L^n\left(\frac{s}{q},\frac{t}{q}\right) = F_L^m\left(\frac{s}{q},\frac{t}{q}\right)$$

for some $n > m$. Since F_L is invertible, this implies

$$F_L^{n-m}\left(\frac{s}{q},\frac{t}{q}\right) = \left(\frac{s}{q},\frac{t}{q}\right).$$

Hence every rational point is periodic. Since rational points are dense in $\mathbb{T}^2$, periodic points are dense.

Conversely, suppose (x,y) is periodic. Then for some n, $F_L^n(x,y) = (x,y)$. Write

$$L^n = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Then $F_L^n(x,y) = (ax+by, cx+dy) \pmod 1$. Thus there exist integers $k,\ell \in \mathbb{Z}$ such that $ax+by = x+k$ and $cx+dy = y+\ell$. Equivalently, $(a-1)x+by = k$ and $cx+(d-1)y = \ell$. Since $L^n - I$ is invertible, we can solve for x and y, obtaining rational numbers. Therefore, periodic points are exactly rational points.

To count $P_n(F_L)$, define $G = F_L^n - \mathrm{id}$. The fixed points of F_L^n are exactly the preimages of $(0,0)$ under G. These correspond to lattice points in the parallelogram

$$(L^n - I)([0,1)\times[0,1)).$$

The number of such points is the absolute value of the determinant $|\det(L^n - I)|$. Since the eigenvalues of L^n are λ_1^n and λ_1^{-n}, we have

$$|\det(L^n - I)| = |(\lambda_1^n - 1)(\lambda_1^{-n} - 1)| = \lambda_1^n + \lambda_1^{-n} - 2.$$

Therefore, $P_n(F_L) = \lambda_1^n + \lambda_1^{-n} - 2$. ∎

Lemma 4.2. The area of a parallelogram with integer vertices is the number of lattice points it contains, where points on the edges are counted as half, and all vertices are counted as a single point.

Proof. Let A denote the area of the parallelogram. Count the lattice points inside the parallelogram in the prescribed way, and denote this number by N.

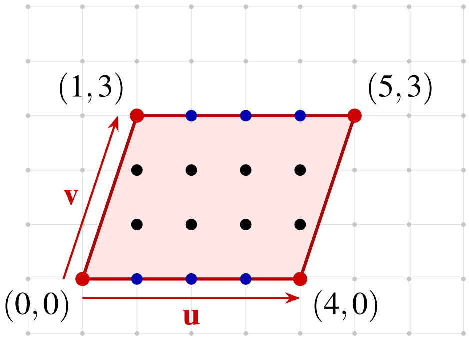

Consider the canonical tiling of the plane by translates of this parallelogram by integer multiples of its edge vectors. Let ℓ be the longest diagonal of the parallelogram.

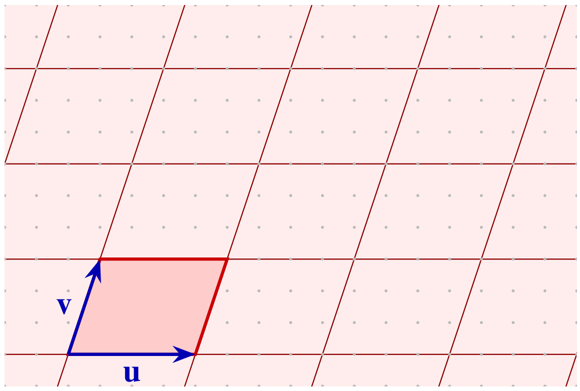

We compare this tiling with the large square $[0,n) \times [0,n)$. The tiles entirely contained in the square cover a slightly smaller square, while the tiles intersecting the square are contained in a slightly larger square.

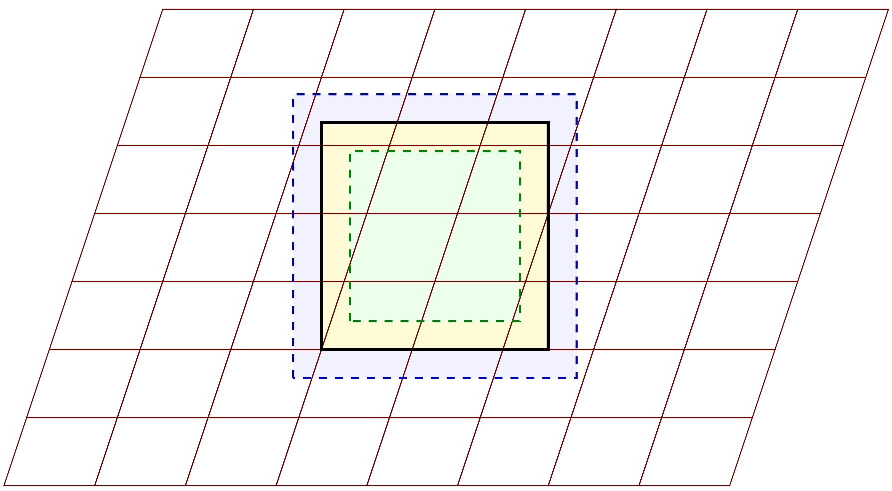

Therefore, for n large, the number of tiles intersecting the square is approximately $\frac{n^2}{A}$. On the other hand, the square contains n^2 integer lattice points. Since each tile contributes N

lattice points in the prescribed counting convention, we get $N \cdot \frac{n^2}{A} \approx n^2$. Letting $n \to \infty$, we obtain $N = A$. Thus the area equals the lattice-point count with the stated boundary convention. □

4.3 Topological Mixing

We now introduce a stronger form of topological transitivity.

Definition 4.2 (topological mixing). A continuous map $f : X \to X$ is said to be topologically mixing if for any two non-empty open sets $U, V \subseteq X$, there exists $N \in \mathbb{N}$ such that $f^n(U) \cap V \neq \emptyset$ for every $n > N$.

Topological mixing means that every open set eventually intersects every other open set under all sufficiently large iterates. Thus mixing is stronger than merely having one dense orbit.

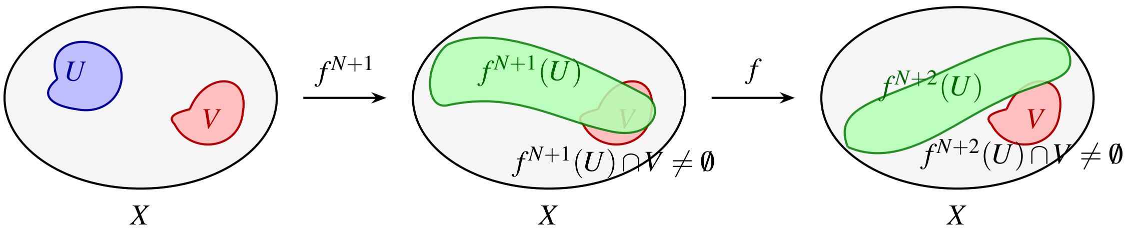

Proposition 4.3. Isometries are not topologically mixing.

Proof. Let $f : X \to X$ be an isometry, so $d(f(x), f(y)) = d(x, y)$ for all $x, y \in X$. Choose three distinct points $x, y, z \in X$. Let $\delta = \frac{1}{4} \min(d(x,y), d(y,z), d(z,x))$. Also, let U, V_1, V_2 be the open δ-balls around x, y, z, respectively.

Since f is an isometry, the diameter of $f^n(U)$ is the same as the diameter of U, and hence is at most 2δ. On the other hand, any point of V_1 and any point of V_2 are more than 2δ apart. Therefore $f^n(U)$ cannot intersect both V_1 and V_2. For each n, at least one of the intersections $f^n(U) \cap V_1$ and $f^n(U) \cap V_2$ must be empty. Hence f is not topologically mixing. □

Proposition 4.4. Expanding maps of S^1 are topologically mixing.

Proof. Let $f : S^1 \to S^1$ be an expanding map such that $|f'(x)| \geq \lambda > 1$ for all $x \in S^1$. Let $F : \mathbb{R} \to \mathbb{R}$ be a lift of f.

If $[a, b] \subseteq \mathbb{R}$ is an interval, then by the mean value theorem, there exists $c \in (a, b)$ such that

$$|F(b) - F(a)| = |F'(c)||b - a| \geq \lambda |b - a|.$$

Therefore the length of any interval is increased by a factor of at least λ under F. After n iterates, the length increases by at least λ^n. Hence for any interval $I \subseteq \mathbb{R}$, there exists

$n \in \mathbb{N}$ such that the length of $F^n(I)$ exceeds 1. Therefore the projection of $F^n(I)$ to S^1 covers the whole circle.

Since every non-empty open set of S^1 contains an interval, it follows that for every non-empty open set $U \subseteq S^1$, some iterate $f^n(U)$ contains S^1. Thus f is topologically mixing. □

We now study topological mixing for the hyperbolic toral automorphism.

Lemma 4.3. Let $T_t^w : \mathbb{T}^n \to \mathbb{T}^n$ be defined by

$$T_t^w(x_1,\dots,x_n) = (x_1 + tw_1,\dots,x_n + tw_n) \pmod 1.$$

Then the flow T_t^w is minimal if and only if the numbers $w_1,\dots,w_n$ are rationally independent, meaning that

$$\sum_{i=1}^{n} k_i w_i \neq 0$$

for any integers $k_1,\dots,k_n$, unless $k_1 = \dots = k_n = 0$.

Proof. Suppose first that $w_1,\dots,w_n$ are rationally independent. It is enough to find a time t_0 such that the map $T_{t_0}^w$ is minimal. By the criterion for toral translations, this is true provided

$$t_0 \sum_{i=1}^{n} k_i w_i$$

is never an integer for any non-zero integer vector $(k_1,\dots,k_n)$.

For each fixed non-zero integer vector and each integer k, there is at most one value of t_0 such that

$$t_0 \sum_{i=1}^{n} k_i w_i = k.$$

There are only countably many such forbidden values of t_0. Hence there exists a choice of t_0 avoiding all of them. Therefore the flow is minimal.

Conversely, suppose that $w_1,\dots,w_n$ are rationally dependent. Then there exists a non-zero integer vector $(k_1,\dots,k_n)$ such that

$$\sum_{i=1}^{n} k_i w_i = 0.$$

Define

$$\phi(x) = \sin\left(2\pi \sum_{i=1}^{n} k_i x_i\right).$$

Then ϕ is continuous, nonconstant, and invariant under the flow T_t^w. Therefore level sets of ϕ give nontrivial closed invariant sets. This contradicts minimality. Hence the flow is minimal only when the w_i's are rationally independent. □

Proposition 4.5. The automorphism $F_L(x,y) = (2x+y, x+y) \pmod 1$ is topologically mixing.

Proof. Let $U, V \subseteq \mathbb{T}^2$ be non-empty open sets. The expanding eigendirection of L has irrational slope. More precisely, the L-invariant family of lines $y = \frac{\sqrt{5}-1}{2}x + \text{constant}$ projects to $\mathbb{T}^2$ as a family of dense orbits of a linear flow with irrational slope $w = \left(1, \frac{\sqrt{5}-1}{2}\right)$. By the previous lemma, this flow is minimal. Hence the projection of each such line is dense in the torus.

Therefore the open set U contains a small segment J of an expanding line. Since F_L expands this direction, the length of $F_L^n(J)$ tends to infinity as $n \to \infty$.

Choose $\varepsilon > 0$ such that V contains an ε-ball. Since sufficiently long segments of irrational-slope lines intersect every ε-ball on the torus, there exists $T > 0$ such that every expanding segment of length at least T intersects V. Since the length of $F_L^n(J)$ grows exponentially, there exists $N \in \mathbb{N}$ such that for every $n \geq N$, the segment $F_L^n(J)$ has length at least T. Therefore $F_L^n(J) \cap V \neq \emptyset$ for every $n \geq N$. Since $J \subseteq U$, we get $F_L^n(U) \cap V \neq \emptyset$ for every $n \geq N$. Hence F_L is topologically mixing. □

4.4 Symbolic Dynamical Systems

One of the most efficient ways to study complicated dynamics is coding. The idea is to divide the phase space into finitely many pieces, and then follow an orbit only by recording which piece it is in at each time. This procedure may lose some information, but it often captures the essential dynamical behaviour. In particular, symbolic dynamics gives a way of converting a geometric dynamical system into a combinatorial one.

We first discuss fractals and Cantor sets. Recall that the Cantor set is one of the most common but unusual objects in analysis. Roughly speaking, it is constructed by repeatedly removing the middle third of each interval.

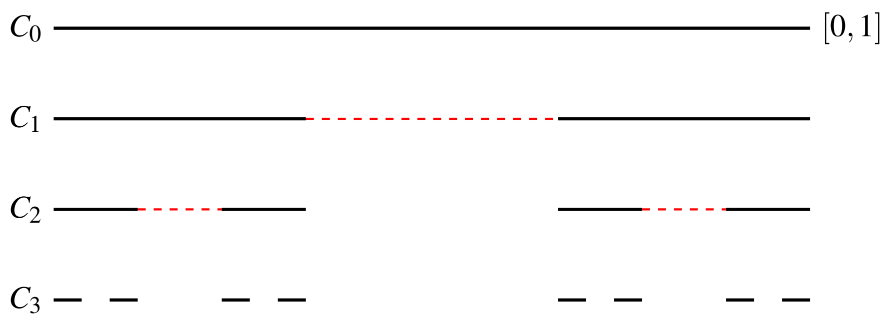

Figure 4.1: Cantor set

Lemma 4.4. The middle-third Cantor set C is the collection of numbers in $[0,1]$ that can be written in ternary expansion without using the digit 1.

Proof. The open middle third $\left(\frac{1}{3}, \frac{2}{3}\right)$ is exactly the set of numbers whose first ternary digit must be 1. These are precisely the numbers that cannot be written with first ternary digit 0 or 2.

At the next step, the removed middle thirds of the remaining intervals are exactly those numbers whose second ternary digit must be 1. Continuing inductively, the points that remain forever are precisely those numbers whose ternary expansions use only the digits 0 and 2. □

Lemma 4.5. The ternary Cantor set C is totally disconnected.

Proof. Let $C_0 = [0,1]$. Let C_1 be the union of the two closed intervals remaining after removing the open middle third: $\left(\frac{1}{3}, \frac{2}{3}\right)$. Thus C_1 consists of two intervals of length 3^{-1}.

Repeating the procedure, C_n consists of 2^n intervals, each of length 3^{-n}, and

$$C = \bigcap_{n=0}^{\infty} C_n.$$

Given any two distinct points of C, they eventually lie in different components of some C_n. These components are separated by open intervals that were removed in the construction. Hence the two points can be separated by disjoint open sets in C. Therefore, C is totally disconnected. □

Lemma 4.6. The ternary Cantor set is uncountable.

Proof. Every point of the Cantor set has a ternary expansion using only 0 and 2. Write such a point as

$$x = \sum_{i=1}^{\infty} \frac{x_i}{3^i} \quad \text{where } x_i \in \{0,2\}.$$

Define

$$f(x) = \sum_{i=1}^{\infty} \frac{x_i/2}{2^i}.$$

Since $x_i/2 \in \{0,1\}$, this gives a binary expansion of a point in $[0,1]$. This map is surjective onto $[0,1]$. Since $[0,1]$ is uncountable, the Cantor set must also be uncountable. □

Definition 4.3 (Cantor set). A set homeomorphic to the ternary Cantor set is called a Cantor set.

Cantor sets are closely related to contractions. For example, the map $f(x) = \frac{x}{3}$ is a contraction on the Cantor set.

Proposition 4.6. There exists a point $x \in S^1$ whose orbit under the doubling map $E_2(x) = 2x \pmod 1$ is dense in S^1.

Proof. We use binary expansions. Every point in S^1 can be represented by a binary expansion $x = 0.x_1x_2x_3\cdots$. The doubling map $E_2(x) = 2x \pmod 1$ acts on the binary expansion by shifting the digits to the left:

$$0.x_1x_2x_3\cdots \mapsto 0.x_2x_3x_4\cdots.$$

Construct an infinite binary sequence by concatenating all binary words of length 1, then all binary words of length 2, then all binary words of length 3, and so on. Let x be the point whose binary expansion is this infinite sequence.

By construction, every finite binary block appears somewhere in the binary expansion of x. Therefore, under iterations of E_2, the orbit of x enters every binary cylinder interval. These cylinder intervals form a basis for the topology of S^1. Hence the orbit of x is dense in S^1. □

The above proof works for the map $E_m(x) = mx \pmod 1$ using expansions in base m.

Proposition 4.7. There exists a point $x \in S^1$ such that the closure of its orbit under $E_3(x) = 3x \pmod 1$ coincides with the standard middle-third Cantor set K. In particular, K is E_3-invariant and contains a dense orbit.

Proof. Recall that the middle-third Cantor set K consists exactly of those points in $[0,1]$ whose ternary expansions use only the digits 0 and 2.

The map $E_3(x) = 3x \pmod 1$ acts on ternary expansions by shifting digits to the left. Therefore K is invariant under E_3. We now construct a point whose orbit is dense in K. Define a map $h : K \to [0,1]$ by changing the ternary digits $0,2$ into binary digits $0,1$. More precisely, if $x = 0.x_1x_2x_3\cdots$ is the ternary expansion of $x \in K$, where each $x_i \in \{0,2\}$, define

$$h(x) = 0.\frac{x_1}{2}\frac{x_2}{2}\frac{x_3}{2}\cdots$$

in binary.

This map is continuous and nondecreasing. It is one-to-one except at points corresponding to binary rationals, which have two binary expansions.

Let $D \subseteq [0,1]$ be a dense set of points that does not contain binary rationals. Then $h^{-1}(D)$ is dense in K. Choose a point in $[0,1]$ whose orbit under E_2 is dense. Pulling it back under h, we obtain a point in K whose orbit under E_3 is dense in K. □

4.5 Sequence Spaces and Shifts

We now give a systematic introduction to symbolic dynamics.

Definition 4.4 (sequence spaces). For each natural number $N \geq 2$, define the space of two-sided sequences of N symbols by

$$\Omega_N = \{w = (\dots, w_{-1}, w_0, w_1, \dots) : w_i \in \{0, 1, \dots, N-1\} \text{ for all } i \in \mathbb{Z}\}.$$

Similarly, define the one-sided sequence space by

$$\Omega_N^R = \{w = (w_0, w_1, w_2, \dots) : w_i \in \{0, 1, \dots, N-1\} \text{ for all } i \in \mathbb{N}_0\}.$$

We think of Ω_N as the space of all bi-infinite symbolic orbits. Each coordinate records the symbol visited at a particular time.

Definition 4.5 (shift map). The two-sided shift map

$$\sigma_N : \Omega_N \to \Omega_N \quad \text{is defined by} \quad (\sigma_N(w))_i = w_{i+1}.$$

Similarly, the one-sided shift map

$$\sigma_N^R : \Omega_N^R \to \Omega_N^R \quad \text{is defined by} \quad (\sigma_N^R(w))_i = w_{i+1}.$$

Thus σ_N moves the sequence one step to the left. The symbol at time $i+1$ becomes the symbol at time i.

Definition 4.6 (cylinder sets). Given a finite block $\alpha = (\alpha_{-m}, \dots, \alpha_m)$, we define the cylinder set

$$C_\alpha^m = \{w \in \Omega_N : w_i = \alpha_i \text{ for } -m \leq i \leq m\}.$$

Thus a cylinder set consists of all sequences whose coordinates in a prescribed finite window are fixed.

Cylinder sets form a basis for the topology on Ω_N. Intuitively, two sequences are close if they agree on a long central block.

Proposition 4.8. Periodic points for the shifts σ_N and σ_N^R are dense in Ω_N and Ω_N^R, respectively. Moreover, both transformations are topologically mixing, and

$$P_n(\sigma_N) = P_n(\sigma_N^R) = N^n.$$

Proof. We prove the result for σ_N. The proof for σ_N^R is similar.

Periodic orbits for the shift are precisely periodic sequences. Indeed, $\sigma_N^m(w) = w$ if and only if $w_{n+m} = w_n$ for all $n \in \mathbb{Z}$.

To prove density of periodic points, it is enough to find a periodic point in every cylinder. Let C_α^m be a cylinder determined by the block $\alpha_{-m}, \dots, \alpha_m$. Repeat this finite block periodically in both directions. That is, define w by requiring $w_n = \alpha_{n'}$, where

$$n' \equiv n \pmod{2m+1} \quad \text{where } -m \leq n' \leq m.$$

Then $w \in C_\alpha^m$, and w is periodic with period $2m+1$. Hence periodic points are dense.

Next, every periodic sequence of period n is determined uniquely by its coordinates $w_0, w_1, \ldots, w_{n-1}$. There are N^n possible choices. Therefore $P_n(\sigma_N) = N^n$.

Finally, we show topological mixing. It is enough to check this on cylinder sets. Let C_α^m and C_β^m be two cylinders. We want to show that for all sufficiently large n, $\sigma_N^n(C_\alpha^m) \cap C_\beta^m \neq \emptyset$. Take $n = 2m+1+k$, where $k > 0$. We can construct a sequence w such that its central block agrees with α, and after shifting n steps, its central block agrees with β. Since the symbols in the gap are arbitrary, such a sequence always exists.

Therefore, $\sigma_N^n(C_\alpha^m) \cap C_\beta^m \neq \emptyset$ for all sufficiently large n. Hence σ_N is topologically mixing. □

> Definition 4.7 (symbolic dynamical system). The restriction of σ_N or σ_N^R to any closed shift-invariant subset of Ω_N or Ω_N^R, respectively, is called a symbolic dynamical system.

4.6 Topological Markov Chains

One of the most important classes of symbolic dynamical systems is given by topological Markov chains.

Let $A = (a_{ij})_{i,j=0}^{N-1}$ be an $N \times N$ matrix whose entries are either 0 or 1. We interpret $a_{ij} = 1$ to mean that the symbol i is allowed to be followed by the symbol j.

Define

$$\Omega_A = \{w \in \Omega_N : a_{w_n w_{n+1}} = 1 \text{ for all } n \in \mathbb{Z}\}.$$

> Definition 4.8 (topological Markov chain). The restriction $\sigma_A = \sigma_N|_{\Omega_A}$ is called the topological Markov chain determined by the matrix A. It is also called a subshift of finite type.

Thus Ω_A consists of all bi-infinite sequences satisfying the transition rules encoded by A.

> Definition 4.9 (transitive matrix). A matrix A is said to be transitive if for every pair of symbols i, j, there exists $n \in \mathbb{N}$ such that $(A^n)_{ij} > 0$. Equivalently, there exists an admissible path from i to j.

Lemma 4.7. If A is transitive and $\alpha = (\alpha_{-k}, \ldots, \alpha_k)$ is admissible, meaning that $a_{\alpha_i \alpha_{i+1}} = 1$ for $i = -k, \ldots, k-1$, then the cylinder $C^k_{\alpha,A} = \Omega_A \cap C^k_\alpha$ is non-empty and contains a periodic point.

Proof. Since A is transitive, there exists $m > 0$ such that there is an admissible path from α_k back to α_{-k}. Therefore we can extend the finite admissible word $\alpha_{-k}, \ldots, \alpha_k$ to a longer admissible word that begins and ends with α_{-k}.

Repeating this longer word periodically gives a periodic sequence in Ω_A. This sequence lies in the cylinder $C^k_{\alpha,A}$. Hence the cylinder is nonempty and contains a periodic point. □

Proposition 4.9. If A is a transitive matrix, then the topological Markov chain σ_A is topologically mixing and its periodic orbits are dense in Ω_A.

Proof. The density of periodic orbits follows from the previous lemma, because every non-empty cylinder contains a periodic point.

We now prove topological mixing. Let $C^k_{\alpha,A}$ and $C^k_{\beta,A}$ be two non-empty cylinders. We need to show that for all sufficiently large n, $\sigma^n_A(C^k_{\alpha,A}) \cap C^k_{\beta,A} \neq \emptyset$.

Since A is transitive, there is an admissible transition from the final symbol of α to the initial symbol of β. Hence, for n sufficiently large, we can insert a finite admissible connecting word between α and β.

This produces an admissible sequence whose first block is α and whose block after n shifts is β. Therefore, $\sigma^n_A(C^k_{\alpha,A}) \cap C^k_{\beta,A} \neq \emptyset$ for all sufficiently large n. Hence σ_A is topologically mixing. □

4.7 Coding of the Toral Automorphism

We now return to the hyperbolic toral automorphism $F(x,y) = (2x+y, x+y) \pmod 1$. The goal is to code the dynamics of F using a symbolic dynamical system.

Theorem 4.1. For the map $F(x,y) = (2x+y, x+y) \pmod 1$ of the torus, there exists a semiconjugacy $h : \Omega_A \to \mathbb{T}^2$ such that

$$F \circ h = h \circ \sigma_5|_{\Omega_A} \quad \text{where} \quad A = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{pmatrix}.$$

Definition 4.10 (Markov partition). A Markov partition is a finite cover $\mathcal{R} = \{R_0, \ldots, R_{m-1}\}$ of $\mathbb{T}^2$ by proper rectangles such that:

(i) the interiors are disjoint: $\mathrm{Int}(R_i) \cap \mathrm{Int}(R_j) = \emptyset$ for $i \neq j$

(ii) whenever $x \in \mathrm{Int}(R_i)$ and $f(x) \in \mathrm{Int}(R_j)$, the local unstable and stable pieces satisfy $W^u_{R_j}(f(x)) \subseteq f(W^u_{R_i}(x))$ and $f(W^s_{R_i}(x)) \subseteq W^s_{R_j}(f(x))$

In Definition 4.10, W^u should be understood as the expanding eigendirection of L, while W^s should be understood as the contracting eigendirection of L. For

$$L = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix},$$

these are the unstable and stable directions of the hyperbolic toral automorphism.

Theorem 4.2. The semiconjugacy between σ_A and F is one-to-one on all periodic points except for the fixed points. Moreover, the number of preimages of any point not negatively asymptotic to the fixed point is bounded.

Proof. We describe the identifications arising from the semiconjugacy, namely the points on the torus that have more than one symbolic preimage.

First, the topological Markov chain σ_A has three fixed points, namely the constant sequences of 0's, 1's, and 4's. However, the toral automorphism F has only one fixed point, namely the origin. Thus these symbolic fixed points are identified under the semiconjugacy.

Next, recall that $P_n(F) = \lambda_1^n + \lambda_1^{-n} - 2$. On the symbolic side, $P_n(\sigma_A) = \mathrm{tr}(A^n)$. By direct computation,

$$\mathrm{tr}(A^n) = \lambda_1^n + \lambda_1^{-n}.$$

Hence, $P_n(\sigma_A) = P_n(F) + 2$. The extra two periodic points correspond precisely to the additional symbolic fixed points that are identified with the origin.

For every point $q \in \mathbb{T}^2$ whose positive and negative iterates avoid the boundaries of the Markov rectangles, the symbolic itinerary is unique. In particular, all periodic points other than the origin belong to this category. Hence the semiconjugacy is one-to-one on all periodic points except for the fixed points.

Points whose iterates lie on the boundaries of the Markov rectangles may have more than one symbolic preimage. These boundary pieces lie along stable and unstable manifolds through the origin. Away from points negatively asymptotic to the fixed point, the number of such possible symbolic itineraries remains bounded. □

Chapter 5

Fractal Dimension and Topological Entropy

Given a topological dynamical system, we would like to know the chaos degree of the system. Two important notions used to measure complexity are fractal dimension and topological entropy. The common root of both notions is the capacity of a set.

Fractal dimension measures how complicated a space is geometrically. Topological entropy measures how complicated the orbits of a dynamical system are. Roughly speaking, topological entropy measures the exponential growth rate of distinguishable orbit segments.

5.1 Capacity

Recall that for a compact metric space, we can define the size or capacity of a set in a similar way to the notion of volume.

Definition 5.1 (capacity). Suppose X is a compact space with metric d. A set $E \subseteq X$ is said to be r-dense if

$$X \subseteq \bigcup_{x \in E} B_d(x,r),$$

where $B_d(x,r)$ is the ball of radius r around x with respect to the metric d. We define the r-capacity of (X,d) to be the minimal cardinality $S_d(r)$ of an r-dense set.

In other words, $S_d(r)$ is the smallest number of balls of radius r needed to cover X.

Example 5.1. Let $X = [0,1]$ with the usual metric. Then $S_d(r)$ is approximately $\frac{1}{2r}$. Indeed, each ball of radius r has length approximately $2r$. Hence it takes roughly $\frac{1}{2r}$ balls of radius r to cover a unit interval.

Example 5.2. Let $X = [0,1]^2$ with the usual metric. Then $S_d(r)$ is approximately $\frac{1}{r^2}$. This agrees with the intuition that a two-dimensional square should require about $\frac{1}{r^2}$ small balls of radius r to cover it.

Example 5.3. Let $X = C$ be the ternary Cantor set with the usual metric. If we use closed balls, then $S_d(3^{-i}) = 2^i$. This is because at the i-th stage of the construction, the Cantor set is covered by 2^i intervals of length 3^{-i}.

If we use open balls instead, then a similar estimate holds. For example, $S_d\left(\left(3-\frac{1}{i}\right)^{-i}\right) = 2^i$. The exact formula depends slightly on whether open or closed balls are used, but the asymptotic behaviour is the important feature.

The capacity $S_d(r)$ is not always easy to compute precisely. However, we can often compute its asymptotic behaviour as $r \to 0$. This leads to the notion of box dimension (Definition 5.2).

Definition 5.2 (box dimension). If X is a compact metric space, then the box dimension of X, if the limit exists, is defined by

$$\operatorname{bdim}(X) = \lim_{r\to 0} -\frac{\log S_d(r)}{\log r}.$$

The idea is that if $S_d(r) \approx r^{-s}$, then $\operatorname{bdim}(X) = s$. Thus the box dimension is the exponent describing how quickly the number of small balls needed to cover X grows as the radius tends to 0.

X $\qquad$ $S_d(r)$ balls

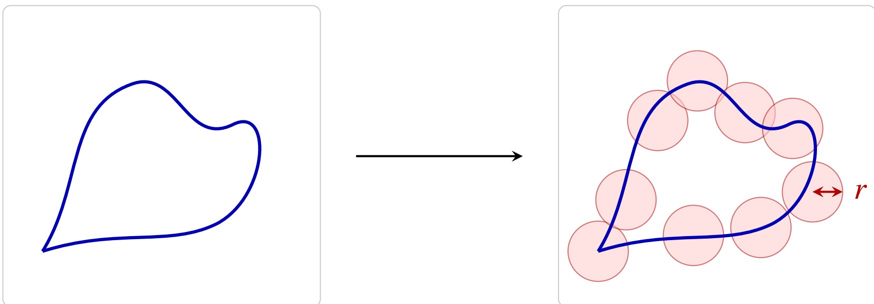

To be completely precise, for a compact metric space X, we should define the lower and upper box dimensions separately.

The lower box dimension is

$$\underline{\operatorname{bdim}}(X) = \liminf_{r\to 0} -\frac{\log S_d(r)}{\log r}.$$

The upper box dimension is

$$\overline{\operatorname{bdim}}(X) = \limsup_{r\to 0} -\frac{\log S_d(r)}{\log r}.$$

We say that the box dimension exists if

$$\underline{\operatorname{bdim}}(X) = \overline{\operatorname{bdim}}(X).$$

In that case, this common value is denoted by $\operatorname{bdim}(X)$.

Recall that a metric space X is totally bounded if for every $\varepsilon > 0$, there exists a finite cover of X by balls of radius ε. The box dimension can also be defined for a totally bounded metric space.

Example 5.4. Let $X = [0,1]$ with the usual metric. Notice that it is not hard to show that

$$\frac{1}{2r} \le S_d(r) \le \frac{1}{2r} + 2.$$

Therefore,

$$-\frac{\log\left(\frac{1}{2r}\right)}{\log r} \le -\frac{\log S_d(r)}{\log r} \le -\frac{\log\left(\frac{1}{2r}+2\right)}{\log r}.$$

By taking $r \to 0$, both the left hand side and the right hand side tend to 1. Hence $\operatorname{bdim}([0,1]) = 1$.

Example 5.5. Let $X = C$ be the ternary Cantor set. At the i^{th} stage of the construction, C is covered by 2^i intervals of length 3^{-i}. Thus $S_d(3^{-i}) = 2^i$. Therefore,

$$\operatorname{bdim}(C) = \lim_{i\to\infty} -\frac{\log 2^i}{\log 3^{-i}} = \lim_{i\to\infty} \frac{i \log 2}{i \log 3} = \frac{\log 2}{\log 3}.$$

Thus

$$\operatorname{bdim}(C) = \frac{\log 2}{\log 3}.$$

Example 5.6. Let C_α be constructed by deleting a middle interval of relative length $1 - \frac{2}{\alpha}$ at each stage. Then the remaining intervals are scaled by a factor of α^{-1} at each step. Hence

$$\operatorname{bdim}(C_\alpha) = \frac{\log 2}{\log \alpha}.$$

This increases to 1 as $\alpha \to 2$, corresponding to deleting smaller and smaller intervals. It decreases to 0 as $\alpha \to \infty$, corresponding to deleting larger and larger intervals.

This shows that box dimension can change under a homeomorphism, since these Cantor sets are pairwise homeomorphic but may have different box dimensions.

Example 5.7. Consider the two-sided sequence space Ω_N with metric d_λ. Recall that cylinders are balls in this metric. More precisely, if $\alpha = \alpha_{1-n} \cdots \alpha_{n-1}$, then

$$C_{\alpha_{1-n}\cdots\alpha_{n-1}} = B_{d_\lambda}(\alpha, \lambda^{1-n}).$$

Hence we need N^{2n-1} balls of radius λ^{1-n} to cover Ω_N. Therefore,

$$\operatorname{bdim}(\Omega_N, d_\lambda) = \lim_{n\to\infty} -\frac{\log N^{2n-1}}{\log \lambda^{1-n}} = \lim_{n\to\infty} \frac{(2n-1)\log N}{(n-1)\log \lambda} = 2\frac{\log N}{\log \lambda}.$$

Thus

$$\operatorname{bdim}(\Omega_N, d_\lambda) = 2\frac{\log N}{\log \lambda}.$$

Similar to the Cantor set example, the box dimension decreases as λ increases, corresponding to the rapid decrease of the radius of cylinders for large λ.

Proposition 5.1. Given a totally bounded metric space (X,d), for any $a>0$,

$$\operatorname{bdim}(X,d)=\operatorname{bdim}(X,ad).$$

Proof. Let ad denote the rescaled metric $(ad)(x,y)=ad(x,y)$. A ball of radius ar with respect to the metric ad is the same as a ball of radius r with respect to the metric d. Hence, $S_{ad}(ar)=S_d(r)$. Therefore,

$$\operatorname{bdim}(X,ad)=\lim_{r\to 0}-\frac{\log S_{ad}(r)}{\log r}.$$

By replacing r with ar, we obtain

$$\operatorname{bdim}(X,ad)=\lim_{r\to 0}-\frac{\log S_{ad}(ar)}{\log (ar)}=\lim_{r\to 0}-\frac{\log S_d(r)}{\log a+\log r}.$$

Since $\frac{\log a}{\log r}\to 0$ as $r\to 0$, we get

$$\operatorname{bdim}(X,ad)=\lim_{r\to 0}-\frac{\log S_d(r)}{\log r}=\operatorname{bdim}(X,d).$$

□

5.2 Topological Entropy

So far, most indicators of complexity that we have met are qualitative. These include topological transitivity, minimality, density of periodic points, chaos, and topological mixing.

The only quantitative measure of complexity we have seen so far is the growth rate of periodic orbits. In this section, we introduce topological entropy, which provides a quantitative way to study the topological complexity of a dynamical system.

Definition 5.3 (Bowen metric). Let $f:X\to X$ be a continuous map of a compact metric space X with metric d. For $n\in\mathbb{N}$, define the metric d_n^f by

$$d_n^f(x,y)=\max_{0\le i\le n-1} d(f^i(x),f^i(y)).$$

Thus $d_n^f(x,y)$ compares the first n iterates of x and y. Two points are close in the metric d_n^f if their orbit segments $x,f(x),\ldots,f^{n-1}(x)$ and $y,f(y),\ldots,f^{n-1}(y)$ remain close for n

steps.

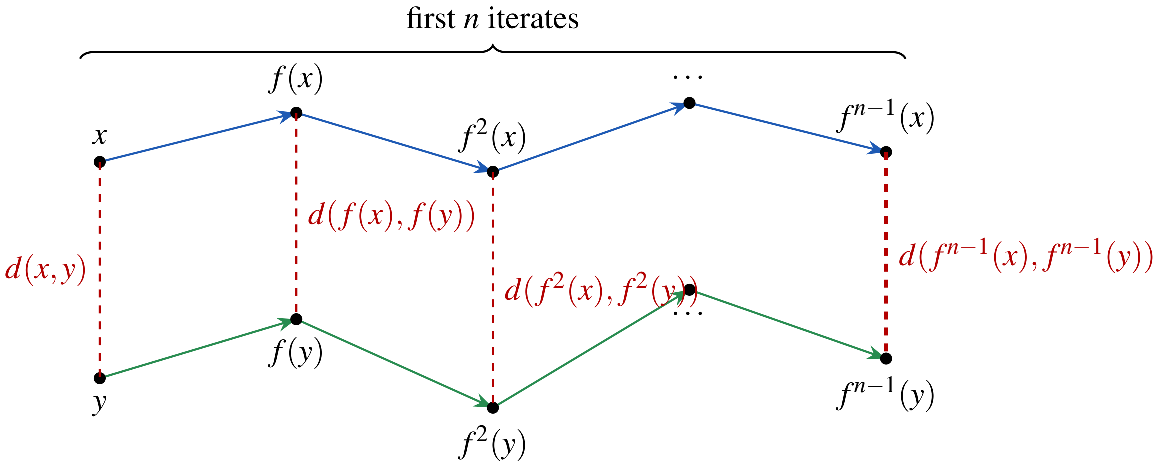

Definition 5.4 ((n,r)-dense set). For $x \in X$, $r > 0$, and $n \in \mathbb{N}$, define the Bowen ball

$$B_f(x,r,n) = \{y \in X : d_n^f(x,y) < r\}.$$

A set $E \subseteq X$ is said to be r-dense with respect to d_n^f, or (n,r)-dense, if

$$X \subseteq \bigcup_{x \in E} B_f(x,r,n).$$

We denote by $S_d(f,r,n)$ the minimal cardinality of an (n,r)-dense set.

Thus $S_d(f,r,n)$ is the smallest number of orbit segments of length n needed to approximate all orbit segments up to accuracy r.

Definition 5.5 (topological entropy). Define

$$h_d(f,r) = \limsup_{n\to\infty} \frac{1}{n} \log S_d(f,r,n).$$

Since $h_d(f,r)$ does not decrease as $r \to 0$, the following limit is well-defined:

$$h_d(f) = \lim_{r\to 0} h_d(f,r).$$

We call $h(f) = h_{\text{top}}(f) = h_d(f)$ the topological entropy of f.

Topological entropy measures the exponential growth rate of distinguishable orbit segments. If a system has many orbit segments that remain distinguishable for a long time, then the entropy is large. If the number of distinguishable orbit segments grows slowly, then the entropy is small.

Proposition 5.2. If d' is a metric on X equivalent to d, then $h_{d'}(f) = h_d(f)$.

Proof. Since d and d' are equivalent metrics on the compact space X, the identity map $\mathrm{id} : (X,d) \to (X,d')$ is a homeomorphism. By compactness, it is uniformly continuous.

Hence, given $r > 0$, there exists $\delta = \delta(r) > 0$ such that if $d'(x_1,x_2) < \delta$, then then $d(x_1,x_2) < r$. Thus every δ-ball in the metric d' is contained in an r-ball in the metric d. The same containment holds for the corresponding Bowen balls. Therefore, for every n,

$$S_{d'}(f,\delta,n) \geq S_d(f,r,n).$$

It follows that $h_{d'}(f) \geq h_d(f)$. By interchanging d and d', we also get $h_d(f) \geq h_{d'}(f)$. The result follows. □

Corollary 5.1. Topological entropy is an invariant of topological conjugacy.

Proof. Let $f : X \to X$ and $g : Y \to Y$ be topologically conjugate via a homeomorphism $h : X \to Y$. Fix a metric d on X, and define a metric d' on Y by

$$d'(y_1,y_2) = d(h^{-1}(y_1),h^{-1}(y_2)).$$

Then h is an isometry from (X,d) to (Y,d'). Hence $h_d(f) = h_{d'}(g)$. Since topological entropy is independent of the choice of equivalent metric, we obtain $h_{\text{top}}(f) = h_{\text{top}}(g)$. □

Proposition 5.3. The topological entropy of contractions and isometries is zero. In particular, any translation T_γ of the torus, or any linear flow T_t^w on the torus, has zero entropy.

Proof. Suppose X is a compact metric space and $f : X \to X$ is 1-Lipschitz. Then $d(f(x), f(y)) \leq d(x,y)$ for all $x,y \in X$.

Hence, for every n,

$$d_n^f(x,y) = \max_{0 \leq i \leq n-1} d(f^i(x), f^i(y)) \leq d(x,y).$$

Thus the number $S_d(f,r,n)$ is bounded above by $S_d(r)$, which does not depend on n. Therefore, $h_{\text{top}}(f) = 0$. Since isometries are 1-Lipschitz, isometries also have zero entropy. □

We can also define topological entropy using covers by sets of small diameter in the Bowen metric.

Let $D_d(f,r,n)$ be the minimal number of sets whose diameter in the metric d_n^f is less than r, and whose union covers X.

Lemma 5.1. For every $r > 0$, the limit

$$\widetilde{h}_d(f,r) = \lim_{n\to\infty} \frac{1}{n} \log D_d(f,r,n)$$

exists.

Proof. Suppose A is a set of d_n^f-diameter less than r, and B is a set of d_m^f-diameter less than r. Then $A \cap f^{-n}(B)$ has d_{n+m}^f-diameter less than r.

Let $\mathcal{U}$ be a cover of X by $D_d(f,r,n)$ sets of d_n^f-diameter less than r, and let $\mathcal{B}$ be a cover of X by $D_d(f,r,m)$ sets of d_m^f-diameter less than r. Then the collection of all sets $A \cap f^{-n}(B)$, where $A \in \mathcal{U}$ and $B \in \mathcal{B}$, covers X. It contains at most $D_d(f,r,n)D_d(f,r,m)$ sets, and each of these sets has d_{n+m}^f-diameter less than r. Therefore,

$$D_d(f,r,n+m) \leq D_d(f,r,n)D_d(f,r,m).$$

Let $a_n = \log D_d(f,r,n)$. Then, $a_{n+m} \leq a_n + a_m$. By the subadditive lemma, the limit

$$\lim_{n\to\infty} \frac{a_n}{n}$$

exists. Hence $\widetilde{h}_d(f,r)$ exists. □

Proposition 5.4. If

$$\underline{h}_d(f,r) = \liminf_{n\to\infty} \frac{1}{n} \log S_d(f,r,n),$$

then

$$\lim_{r\to 0} \widetilde{h}_d(f,r) = \lim_{r\to 0} \underline{h}_d(f,r) = \lim_{r\to 0} h_d(f,r) = h_{\text{top}}(f).$$

Proof. The diameter of an r-ball is at most $2r$. Hence every covering by r-balls is also a covering by sets of diameter less than $2r$. Therefore,

$$D_d(f,2r,n) \leq S_d(f,r,n).$$

On the other hand, any set of diameter less than r is contained in an r-ball around each of its points. Hence

$$S_d(f,r,n) \leq D_d(f,r,n).$$

Thus

$$\widetilde{h}_d(f,2r) \leq \underline{h}_d(f,r) \leq h_d(f,r) \leq \widetilde{h}_d(f,r).$$

Taking $r \to 0$, all three expressions give the same value, namely $h_{\text{top}}(f)$. □

Another useful definition of topological entropy is given by separated sets.

Definition 5.6 ((n,r)-separated set). A set $E \subseteq X$ is called (n,r)-separated if for any two distinct points $x,y \in E$, $d_n^f(x,y) \geq r$. Equivalently, for any two distinct points $x,y \in E$, there exists $0 \leq i \leq n-1$ such that

$$d(f^i(x), f^i(y)) \geq r.$$

Let $N_d(f,r,n)$ be the maximal cardinality of an (n,r)-separated set.

Proposition 5.5. We have

$$h_{\text{top}}(f) = \lim_{r\to 0} \limsup_{n\to\infty} \frac{1}{n} \log N_d(f,r,n).$$

Moreover,

$$h_{\text{top}}(f) = \lim_{r\to 0} \liminf_{n\to\infty} \frac{1}{n} \log N_d(f,r,n).$$

Proof. A maximal (n,r)-separated set is (n,r)-dense. Indeed, if it were not (n,r)-dense, then there would be some point outside all r-Bowen balls around the separated set. We could then add this point to the separated set, contradicting maximality. Hence

$$S_d(f,r,n) \le N_d(f,r,n).$$

On the other hand, no Bowen ball of radius $\frac{r}{2}$ can contain two points that are r-apart in the Bowen metric. Therefore,

$$N_d(f,r,n) \le S_d\left(f,\frac{r}{2},n\right).$$

Taking logarithms, dividing by n, taking the appropriate limits, and then letting $r \to 0$, we obtain the desired formula for topological entropy. □

We now discuss some properties of topological entropy.

Proposition 5.6. Topological entropy satisfies the following properties.

(i) If Λ is a closed f-invariant set, then $h_{\text{top}}(f|_\Lambda) \le h_{\text{top}}(f)$

(ii) If

$$X = \bigcup_{i=1}^{m} \Lambda_i,$$

where each Λ_i is closed and f-invariant, then

$$h_{\text{top}}(f) = \max_{1\le i\le m} h_{\text{top}}(f|_{\Lambda_i}).$$

(iii) For $m \in \mathbb{Z}$, $h_{\text{top}}(f^m) = |m| h_{\text{top}}(f)$

(iv) If g is a factor of f, then $h_{\text{top}}(g) \le h_{\text{top}}(f)$

Proof. We explain the ideas behind each statement.

For **(i)**, every orbit segment in Λ is also an orbit segment in X. Hence the number of distinguishable orbit segments in Λ cannot be larger than the number of distinguishable orbit segments in X. The result follows.

For **(ii)**, the space X is the union of finitely many closed invariant pieces. Since the number of Bowen balls needed to cover X is controlled by the sum of the numbers needed to cover the pieces Λ_i, the exponential growth rate is the maximum of the exponential growth rates on the pieces. Hence

$$h_{\text{top}}(f) = \max_{1 \le i \le m} h_{\text{top}}(f|_{\Lambda_i}).$$

For **(iii)**, an orbit segment of length n for f^m corresponds to an orbit segment of length approximately mn for f. Therefore the exponential growth rate is multiplied by $|m|$.

For **(iv)**, suppose g is a factor of f. Then there exists a continuous surjective map $h : X \to Y$ such that $h \circ f = g \circ h$. Every orbit segment of g is the image of an orbit segment of f. Therefore the factor system cannot have more orbit complexity than the original system. The result follows. □

We then discuss some examples of topological entropy.

Example 5.8 (full shift). Consider the full shift $\sigma_N : \Omega_N \to \Omega_N$. We claim that $h_{\text{top}}(\sigma_N) = \log N$. Indeed, an orbit segment of length n is determined by a word of length n, and there are N^n such words. Therefore,

$$h_{\text{top}}(\sigma_N) = \lim_{n\to\infty} \frac{1}{n} \log N^n = \log N.$$

The same computation applies to the one-sided full shift.

Example 5.9 (topological Markov chain). Let A be an $N \times N$ matrix with entries 0 or 1, and let $\sigma_A : \Omega_A \to \Omega_A$ be the corresponding topological Markov chain.

The number of admissible words of length n is governed by the powers of A. More precisely, the number of admissible paths of length n from symbol i to symbol j is $(A^n)_{ij}$. Hence the total number of admissible words grows asymptotically like the spectral radius of A. Therefore, $h_{\text{top}}(\sigma_A) = \log \rho(A)$ where $\rho(A)$ denotes the spectral radius of A.

Example 5.10 (toral automorphism). Consider the hyperbolic toral automorphism induced by

$$L = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

The eigenvalues are $\lambda_1 = \frac{3+\sqrt{5}}{2} > 1$ and $\lambda_1^{-1} = \frac{3-\sqrt{5}}{2} < 1$. The unstable direction is expanded by the factor λ_1, while the stable direction is contracted by λ_1^{-1}. Thus the exponential orbit complexity is controlled by the expanding eigenvalue. Hence, $h_{\text{top}}(F_L) = \log \lambda_1$.

Example 5.11 (circle maps and degree). Let $f : S^1 \to S^1$, and let $F : \mathbb{R} \to \mathbb{R}$ be a lift satisfying $F(x+1) = F(x) + D$, where $D = \deg f$. We will show that

$$h_{\text{top}}(f) \ge \log|\deg(f)|.$$

Thus the degree of a circle map gives a lower bound for its entropy.

First, set $d = |D|$. If $d \leq 1$, then $\log d \leq 0$. Since topological entropy is always non-negative, the inequality is trivial. Hence we may assume that $d \geq 2$. Let $F : \mathbb{R} \to \mathbb{R}$ be a lift of f, so that $F(x+1) = F(x) + D$. Then,

$$F^m(x+1) = F^m(x) + D^m \quad \text{for all } m \geq 0.$$

Choose $\varepsilon \in \left(0, \frac{1}{4}\right)$ so small that

$$|a-b| < \varepsilon \quad \text{implies} \quad |F(a) - F(b)| < \frac{1}{4}.$$

By the intermediate value theorem,

$$|F^{n-1}(1) - F^{n-1}(0)| = d^{n-1}.$$

Thus the image $F^{n-1}([0,1])$ covers an interval of length d^{n-1}. Hence we can choose points $E_n = \{z_0, z_1, \ldots, z_{d^{n-1}-1}\} \subseteq [0,1)$ such that

$$F^{n-1}(z_j) = F^{n-1}(0) + j \cdot \operatorname{sgn}(D)^{n-1}.$$

We claim that the projections of the points in E_n to S^1 form an (n, ε)-separated set.

Suppose, for contradiction, that $z_i \neq z_j$ are not (n, ε)-separated. Then for every $m = 0, \ldots, n-1$, the points $f^m(z_i)$ and $f^m(z_j)$ are within distance ε on the circle. Therefore, for each m, there exists a unique integer $k_m \in \mathbb{Z}$ such that

$$|F^m(z_i) - F^m(z_j) - k_m| < \varepsilon.$$

Write

$$F^m(z_i) = F^m(z_j) + k_m + \delta_m \quad \text{where } |\delta_m| < \varepsilon.$$

Applying F, and using $F(x+k) = F(x) + Dk$, we obtain

$$F^{m+1}(z_i) = F(F^m(z_j) + k_m + \delta_m) = F(F^m(z_j) + \delta_m) + Dk_m.$$

Hence

$$F^{m+1}(z_i) - F^{m+1}(z_j) = F(F^m(z_j) + \delta_m) - F(F^m(z_j)) + Dk_m.$$

By the choice of ε,

$$|F(F^m(z_j) + \delta_m) - F(F^m(z_j))| < \frac{1}{4}.$$

Thus $F^{m+1}(z_i) - F^{m+1}(z_j)$ is within $\frac{1}{4}$ of the integer Dk_m. But by definition, it is also within $\varepsilon < \frac{1}{4}$ of the integer k_{m+1}. Hence

$$|k_{m+1} - Dk_m| < \frac{1}{2}.$$

Since both k_{m+1} and Dk_m are integers, this implies $k_{m+1} = Dk_m$. By induction, $k_{n-1} = D^{n-1}k_0$. Now evaluate the inequality at $m = 0$. Since $z_i, z_j \in [0,1)$, and

$$|z_i - z_j - k_0| < \varepsilon < \frac{1}{4},$$

we must have $k_0 \in \{0, \pm 1\}$. At $m = n-1$, our construction gives

$$F^{n-1}(z_i) - F^{n-1}(z_j) = (i-j)\,\mathrm{sgn}(D)^{n-1}.$$

Using the inequality for $m = n-1$, we get

$$\left|(i-j)\,\mathrm{sgn}(D)^{n-1} - D^{n-1}k_0\right| < \varepsilon < \frac{1}{4}.$$

Equivalently,

$$\left|i - j - d^{n-1}\,\mathrm{sgn}(k_0)\right| < \frac{1}{4}.$$

But $i, j \in \{0, 1, \ldots, d^{n-1} - 1\}$. Hence, $|i-j| < d^{n-1}$. If $k_0 \neq 0$, then $|i - j \pm d^{n-1}| \geq 1$ which contradicts the previous inequality. Therefore, $k_0 = 0$. Then the same inequality forces $i = j$, contradicting the assumption that the two points were distinct.

Hence all d^{n-1} points in E_n are (n, ε)-separated. Therefore, $N_d(f, \varepsilon, n) \geq d^{n-1}$. Taking the limit gives

$$h_{\mathrm{top}}(f) \geq \limsup_{n\to\infty} \frac{1}{n} \log(d^{n-1}) = \log d = \log|\deg(f)|.$$

This result shows that a continuous circle map with degree of absolute value at least 2 must have positive entropy. Intuitively, such a map wraps the circle around itself many times, creating exponentially many distinguishable orbit segments.

Chapter 6

Hausdorff Measure and Hausdorff Dimension

In Chapter 5, we introduced box dimension as a way to measure the size or fractal complexity of a compact metric space. Box dimension is convenient and easy to compute in many examples. However, it has several mathematical shortcomings.

The main problem is that box dimension uses finite covers by balls of the same radius. This makes it less flexible than a measure-theoretic notion of dimension. In this chapter, we introduce Hausdorff measure and Hausdorff dimension, which are more robust and better suited for rigorous analysis in dynamical systems.

Recall that the box dimension of a compact metric space X is defined using the growth rate of the minimal number of balls of radius r needed to cover X. This gives a useful notion of size, but it can behave badly for certain sets.

Example 6.1. Let $X = \mathbb{Q} \cap [0,1]$. Since $\mathbb{Q}$ is dense in $[0,1]$, any covering of X by balls of radius r essentially needs the same number of balls as a covering of the whole interval $[0,1]$. Therefore,

$$\mathrm{bdim}(\mathbb{Q} \cap [0,1]) = \mathrm{bdim}([0,1]) = 1.$$

However, $\mathbb{Q} \cap [0,1]$ is countable. Intuitively, a countable set of points should have dimension 0.

This example shows that box dimension does not have the property of countable stability. That is, even if a set is a countable union of very small pieces, box dimension may still assign it a large dimension.

The root of this problem is that box dimension relies on finite covers by sets of exactly the same scale. To fix this, we introduce Hausdorff dimension, which allows covers by sets of varying and arbitrarily small diameters.

6.1 Hausdorff Measure

Let (X,d) be a metric space and let $E \subseteq X$. If $U \subseteq X$ is a nonempty subset, we define its diameter by

$$|U| = \sup\{d(x,y) : x,y \in U\}.$$

Definition 6.1 (δ-cover). Let $E \subseteq X$, $s \geq 0$, and $\delta > 0$. A δ-cover of E is a countable or finite collection of sets $\{U_i\}_{i=1}^{\infty}$ such that

$$E \subseteq \bigcup_{i=1}^{\infty} U_i \quad \text{and} \quad |U_i| < \delta \text{ for all } i.$$

Thus a δ-cover is a cover of E by sets whose diameters are all smaller than δ. Unlike box dimension, the sets in the cover do not need to have the same size.

Definition 6.2 (Hausdorff premeasure). Let $E \subseteq X$, $s \geq 0$, and $\delta > 0$. We define

$$\mathcal{H}^s_\delta(E) = \inf\left\{\sum_{i=1}^{\infty} |U_i|^s : \{U_i\}_{i=1}^{\infty} \text{ is a } \delta\text{-cover of } E\right\}.$$

As δ decreases, the class of permissible covers becomes smaller. Hence the infimum is taken over fewer possible covers. Therefore,

$$\mathcal{H}^s_\delta(E)$$

is monotonically increasing as $\delta \to 0$. This guarantees that the following limit exists, although it may be infinite.

Definition 6.3 (Hausdorff outer measure). The s-dimensional Hausdorff outer measure of E is defined by

$$\mathcal{H}^s(E) = \lim_{\delta \to 0} \mathcal{H}^s_\delta(E).$$

Equivalently,

$$\mathcal{H}^s(E) = \sup_{\delta > 0} \mathcal{H}^s_\delta(E).$$

Remark 6.1. It is a standard result in geometric measure theory that $\mathcal{H}^s$ is a Borel regular measure.

When $s = 0$, $\mathcal{H}^0(E)$ is the counting measure of E. In other words, it counts the number of points in E.

When $s = 1$ and $E \subseteq \mathbb{R}$, $\mathcal{H}^1(E)$ coincides with the usual Lebesgue outer measure, namely length.

When $s = 2$ and $E \subseteq \mathbb{R}^2$, $\mathcal{H}^2(E)$ is proportional to the usual two-dimensional Lebesgue measure, namely area. The proportionality constant depends on the convention used for diameter.

6.1.1 Hausdorff Dimension

The fundamental property of Hausdorff measure is how it behaves as the parameter s changes. For a fixed set E, the quantity $\mathcal{H}^s(E)$ jumps from ∞ to 0 at a critical value of s. This critical value is the Hausdorff dimension.

Lemma 6.1. Let $E \subseteq X$. If

$$\mathcal{H}^s(E) < \infty,$$

then for any $t > s$,

$$\mathcal{H}^t(E) = 0.$$

Conversely, if

$$\mathcal{H}^t(E) > 0,$$

then for any $s < t$,

$$\mathcal{H}^s(E) = \infty.$$

Proof. Suppose $\{U_i\}_{i=1}^{\infty}$ is a δ-cover of E. Since

$$|U_i| < \delta$$

and $t > s$, we have

$$|U_i|^t = |U_i|^{t-s}|U_i|^s \leq \delta^{t-s}|U_i|^s.$$

Taking the sum over all i, we get

$$\sum_{i=1}^{\infty} |U_i|^t \leq \delta^{t-s} \sum_{i=1}^{\infty} |U_i|^s.$$

Taking the infimum over all δ-covers gives

$$\mathcal{H}^t_\delta(E) \leq \delta^{t-s}\mathcal{H}^s_\delta(E).$$

Now let $\delta \to 0$. Since $t - s > 0$, we have

$$\delta^{t-s} \to 0.$$

If $\mathcal{H}^s(E) < \infty$, then the right hand side tends to 0. Hence

$$\mathcal{H}^t(E) = 0.$$

The converse follows by taking the contrapositive. □

This lemma shows that the graph of $\mathcal{H}^s(E)$ as a function of s has a sharp transition. Before the critical value, the measure is infinite. After the critical value, the measure is zero.

Definition 6.4 (Hausdorff dimension). The Hausdorff dimension of a set $E \subseteq X$ is the critical value

$$\dim_H(E) = \inf\{s \geq 0 : \mathcal{H}^s(E) = 0\}.$$

Equivalently,

$$\dim_H(E) = \sup\{s \geq 0 : \mathcal{H}^s(E) = \infty\}.$$

Proposition 6.1. Hausdorff dimension satisfies the following properties.

(i) Monotonicity: If $E \subseteq F$, then

$$\dim_H(E) \leq \dim_H(F).$$

(ii) Countable stability: If

$$E = \bigcup_{i=1}^{\infty} E_i,$$

then

$$\dim_H(E) = \sup_i \dim_H(E_i).$$

(iii) Comparison with box dimension: For any compact metric space E,

$$\dim_H(E) \leq \underline{\text{bdim}}(E) \leq \overline{\text{bdim}}(E).$$

Proof. We prove each statement separately.

For **(i)**, suppose $E \subseteq F$. Any cover of F is automatically a cover of E. Hence

$$\mathcal{H}^s_\delta(E) \leq \mathcal{H}^s_\delta(F)$$

for all $s \geq 0$ and all $\delta > 0$. Taking $\delta \to 0$, we obtain

$$\mathcal{H}^s(E) \leq \mathcal{H}^s(F).$$

Therefore, if $\mathcal{H}^s(F) = 0$, then $\mathcal{H}^s(E) = 0$. Hence

$$\dim_H(E) \leq \dim_H(F).$$

For **(ii)**, since

$$E_i \subseteq E$$

for every i, monotonicity gives

$$\dim_H(E_i) \leq \dim_H(E)$$

for all i. Hence

$$\sup_i \dim_H(E_i) \leq \dim_H(E).$$

We now prove the reverse inequality.

Let

$$s > \sup_i \dim_H(E_i).$$

Then for every i,

$$s > \dim_H(E_i),$$

so by the definition of Hausdorff dimension,

$$\mathcal{H}^s(E_i) = 0.$$

By countable subadditivity of Hausdorff measure,

$$\mathcal{H}^s(E) = \mathcal{H}^s\left(\bigcup_{i=1}^{\infty} E_i\right) \leq \sum_{i=1}^{\infty} \mathcal{H}^s(E_i) = 0.$$

Thus

$$\mathcal{H}^s(E) = 0.$$

Since this holds for every

$$s > \sup_i \dim_H(E_i),$$

we obtain

$$\dim_H(E) \leq \sup_i \dim_H(E_i).$$

Combining both inequalities gives

$$\dim_H(E) = \sup_i \dim_H(E_i).$$

For **(iii)**, it is clear from the definitions that

$$\underline{\text{bdim}}(E) \leq \overline{\text{bdim}}(E).$$

It remains to prove that

$$\dim_H(E) \leq \underline{\text{bdim}}(E).$$

Let $N(r)$ be the minimal number of sets of diameter at most r needed to cover E. Recall that

$$\underline{\text{bdim}}(E) = \liminf_{r \to 0} \frac{\log N(r)}{-\log r}.$$

Choose

$$s > \underline{\text{bdim}}(E).$$

Then we can find t such that

$$s > t > \underline{\text{bdim}}(E).$$

By the definition of the limit inferior, there exists a sequence $r_k \to 0$ such that

$$\frac{\log N(r_k)}{-\log r_k} < t$$

for all k. This implies

$$\log N(r_k) < -t \log r_k,$$

and hence

$$N(r_k) < r_k^{-t}.$$

For each k, choose a cover

$$\{U_{k,i}\}_{i=1}^{N(r_k)}$$

of E such that

$$|U_{k,i}| \leq r_k.$$

Then this is an r_k-cover of E, and hence

$$\begin{aligned}\mathcal{H}^s_{r_k}(E) &\leq \sum_{i=1}^{N(r_k)} |U_{k,i}|^s \\ &\leq N(r_k) r_k^s \\ &< r_k^{-t} r_k^s \\ &= r_k^{s-t}.\end{aligned}$$

Since $s > t$ and $r_k \to 0$, we get

$$r_k^{s-t} \to 0.$$

Therefore,

$$\mathcal{H}^s(E) = 0.$$

Since this holds for every

$$s > \underline{\text{bdim}}(E),$$

we conclude that

$$\dim_H(E) \leq \underline{\text{bdim}}(E).$$

□

Remark 6.2. Countable stability immediately resolves the problem with rational numbers. Since every single point x satisfies

$$\dim_H(\{x\}) = 0,$$

and since

$$\mathbb{Q} \cap [0,1]$$

is a countable union of points, we get

$$\dim_H(\mathbb{Q} \cap [0,1]) = 0.$$

This agrees with the intuition that a countable set should have dimension zero.

6.1.2 The Mass Distribution Principle

Hausdorff dimension is more robust than box dimension, but it is also harder to compute directly.

Establishing an upper bound for Hausdorff dimension is usually not too difficult. We only need to find one efficient δ-cover. Establishing a lower bound is harder, because we need to show that no cover can be too efficient.

To obtain lower bounds, we use measure theory. The idea is simple: if we can distribute a probability mass on E in such a way that the mass is not too concentrated anywhere, then E must be geometrically large.

Theorem 6.1 (mass distribution principle). Let $E \subseteq X$ be a Borel set. Suppose there exists a probability measure μ supported on E, meaning that

$$\mu(E) = 1 \qquad \text{and} \qquad \mu(X \setminus E) = 0.$$

Suppose there exist constants $c > 0$ and $s \geq 0$ such that for every set $U \subseteq X$,

$$\mu(U) \leq c|U|^s.$$

Then

$$\mathcal{H}^s(E) \geq \frac{1}{c} > 0.$$

Consequently,

$$\dim_H(E) \geq s.$$

Proof. Let

$$\{U_i\}_{i=1}^{\infty}$$

be any arbitrary δ-cover of E. Since μ is supported on E, we have

$$1 = \mu(E).$$

Since

$$E \subseteq \bigcup_{i=1}^{\infty} U_i,$$

we get

$$1 = \mu(E) \leq \mu\left(\bigcup_{i=1}^{\infty} U_i\right) \leq \sum_{i=1}^{\infty} \mu(U_i).$$

Using the assumption

$$\mu(U_i) \leq c|U_i|^s,$$

we obtain

$$1 \leq \sum_{i=1}^{\infty} c|U_i|^s = c\sum_{i=1}^{\infty} |U_i|^s.$$

Dividing by c, we obtain

$$\sum_{i=1}^{\infty} |U_i|^s \geq \frac{1}{c}.$$

Since this holds for every δ-cover of E, taking the infimum over all such covers gives

$$\mathcal{H}^s_\delta(E) \geq \frac{1}{c}.$$

Letting $\delta \to 0$, we get

$$\mathcal{H}^s(E) \geq \frac{1}{c} > 0.$$

By the definition of Hausdorff dimension, this implies

$$\dim_H(E) \geq s.$$

□

6.1.3 Dimension of the Cantor Set

We now apply the preceding tools to compute rigorously the Hausdorff dimension of the standard middle-third Cantor set C. Recall that C is obtained by repeatedly removing the open middle third of every interval.

Theorem 6.2. The Hausdorff dimension of the middle-third Cantor set C is

$$\dim_H(C) = \frac{\log 2}{\log 3}.$$

Proof. Let $s = \frac{\log 2}{\log 3}$. We prove the theorem by showing both inequalities $\dim_H(C) \leq s$ and $\dim_H(C) \geq s$.

First, we prove the upper bound. Recall that

$$C = \bigcap_{k=0}^{\infty} C_k,$$

where C_k consists of 2^k intervals, each of length 3^{-k}. These 2^k intervals form a cover of C. Taking $\delta = 3^{-k}$, we obtain a δ-cover of C. Therefore,

$$\mathcal{H}^s_\delta(C) \leq \sum_{i=1}^{2^k} |U_i|^s.$$

Since each interval has diameter 3^{-k}, we get

$$\sum_{i=1}^{2^k} |U_i|^s = 2^k (3^{-k})^s = 2^k (3^{-k})^{\frac{\log 2}{\log 3}} = 2^k \left(3^{\frac{\log 2}{\log 3}}\right)^{-k} = 2^k (2)^{-k} = 1.$$

Hence, $\mathcal{H}^s(C) \leq 1 < \infty$ so Therefore, $\dim_H(C) \leq s$.

Next, we prove the lower bound. We construct a probability measure μ on C. Distribute total mass 1 uniformly across the Cantor intervals. At the first step, assign mass $\frac{1}{2}$ to each of the two intervals of length $\frac{1}{3}$. At step k, each of the 2^k intervals of length 3^{-k} receives mass 2^{-k}. This defines a probability measure μ supported entirely on C.

Let U be an arbitrary open interval with $|U| < 1$. Choose an integer $k \geq 0$ such that

$$3^{-(k+1)} \leq |U| < 3^{-k}.$$

Since the interval U has length less than 3^{-k}, it can intersect at most one basic interval of C_k. Therefore, the measure of U is bounded by the mass of a single k^{th} level interval: $\mu(U) \leq 2^{-k}$. Now

$$2^{-k} = (3^{-k})^{\frac{\log 2}{\log 3}} = (3^{-k})^s.$$

Also,

$$3^{-k} = 3 \cdot 3^{-(k+1)} \leq 3|U|.$$

Hence

$$2^{-k} = (3^{-k})^s \le (3|U|)^s = 3^s|U|^s.$$

Since $s = \frac{\log 2}{\log 3}$, we have $3^s = 2$ so $\mu(U) \le 2|U|^s$. By the mass distribution principle, taking $c = 2$, we obtain $\mathcal{H}^s(C) \ge \frac{1}{2} > 0$. Therefore, $\dim_H(C) \ge s$. Combining the upper and lower bounds gives the result. □

Remark 6.3. This agrees with our earlier box dimension calculation. However, the Hausdorff dimension computation is stronger, because it is based on Hausdorff measure and uses countable covers of arbitrary small diameters.

6.1.4 Iterated Function Systems and the Moran Equation

The Cantor set is a classical example of an invariant set arising from an iterated function system. This idea is deeply connected with dynamical systems, because many fractal invariant sets arise as attractors or repellers of simple maps.

Definition 6.5 (iterated function system). An iterated function system, or IFS, is a finite collection of contraction maps

$$\{S_1, S_2, \dots, S_m\}$$

on a complete metric space (X, d). Each S_i has a contraction ratio

$$0 < c_i < 1.$$

That is,

$$d(S_i(x), S_i(y)) \le c_i d(x, y)$$

for all $x, y \in X$.

By Hutchinson's theorem, there exists a unique nonempty compact invariant set

$$F \subseteq X$$

such that

$$F = \bigcup_{i=1}^{m} S_i(F).$$

This set F is called the invariant set or attractor of the iterated function system.

Definition 6.6 (open set condition). An iterated function system $\{S_1, \dots, S_m\}$ is said to satisfy the open set condition if there exists a nonempty bounded open set $V \subseteq X$ such that

$$\bigcup_{i=1}^{m} S_i(V) \subseteq V$$

and

$$S_i(V) \cap S_j(V) = \emptyset$$

whenever $i \ne j$.

The open set condition says that the different pieces of the invariant set are sufficiently separated. When this condition holds, the Hausdorff dimension can be computed algebraically.

Theorem 6.3 (Moran equation). Suppose an iterated function system $\{S_1, \dots, S_m\}$ has contraction ratios $c_1, \dots, c_m$ and satisfies the open set condition. Then the Hausdorff dimension and box dimension of its invariant set F coincide, and they are equal to the unique positive real number s satisfying

$$\sum_{i=1}^{m} c_i^s = 1.$$

That is, $\dim_H(F) = \mathrm{bdim}(F) = s$.

The Moran equation abstracts the Cantor set calculation. The middle-third Cantor set is the invariant set of the two contractions $S_1(x) = \frac{x}{3}$ and $S_2(x) = \frac{x}{3} + \frac{2}{3}$. Both maps have contraction ratio $c_1 = c_2 = \frac{1}{3}$. The open set condition is satisfied with $V = (0,1)$. Therefore, the Moran equation becomes $\left(\frac{1}{3}\right)^s + \left(\frac{1}{3}\right)^s = 1$, or equivalently, $3^s = 2$. Hence, Equivalently, $s = \frac{\log 2}{\log 3}$. This recovers the Hausdorff dimension of the middle-third Cantor set.

In dynamical systems, such invariant sets often appear as non-wandering sets or repellers of expanding maps. For instance, certain points whose orbits never escape under an expanding map may form a Cantor-like set. The Hausdorff dimension then measures the geometric size of this surviving invariant set.

Bibliography

[1] Lawrence C. Evans, *Partial Differential Equations*, 2nd ed., Graduate Studies in Mathematics, Vol. 19, American Mathematical Society, Providence, RI, 2010. ISBN: 978-0-8218-4974-3.

[2] Walter A. Strauss, *Partial Differential Equations: An Introduction*, 2nd ed., John Wiley & Sons, Hoboken, NJ, 2008. ISBN: 978-0-470-05456-7.